



2007 ANNUAL REPORT

COMPUTATION DIRECTORATE



LAWRENCE LIVERMORE NATIONAL LABORATORY

LLNL Computation Directorate Annual Report 2007



About the Cover

The main control room of the National Ignition Facility's (NIF's) Integrated Computer Control System (ICCS) is shown on the front cover. ICCS is the most complex real-time control system ever designed for a scientific machine. The simulation in the foreground was created on the Atlas supercomputer and shows a laser beam propagating through the ionized gas inside a NIF hohlraum. The Purple supercomputer shown on the back cover provides more than 100 trillion floating-point operations per second of computing power to help achieve important and diverse science-based simulation missions. Shown in the foreground is a simulation, run on Purple, that helped guide NIF's beryllium capsule designs. The optimized design minimizes penetration of the beryllium into the fusion fuel, thus maximizing fusion burn. (Cover design by Daniel S. Moore)

Production Team

LLNL Associate Director for Computation: Dona L. Crawford

Deputy Associate Directors: Ronald F. Schmucker (Principal Deputy),
Patricia A. Damkroger, and
Michel G. McCoy

Scientific Editors: Van Emden Henson and
James A. Guse

Art Director: Daniel S. Moore

Production Editor: Deanna K. Willis

Editor and Proofreader: Carolin Middleton

Photographers: Lee E. Baker and
Nancy A. Rutter

Print Production: Charlie M. Arteago, Jr., and
Monarch Print Copy and Design Solutions



This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

© 2008. Lawrence Livermore National Security, LLC. All rights reserved. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344.

Contents

Section 1: A Message from the Associate Director

1.00 — Extraordinary Computational Achievements Prevail During a Year of Transition	2
1.01 — Awards and Recognition	5

Section 2: High-Performance Computing

2.00 — Computation Develops Strategies for Preeminent National Security Computing	7
2.01 — Expanded BlueGene/L Brings Predictive, Fully Integrated Calculations Closer to Reality	9
2.02 — BlueGene/L Enables Large-Scale Science Runs	10
2.03 — Next-Generation Development Environment Tools Pave the Way to Sequoia-Scale Computing	12
2.04 — Common Capacity Computing Lowers Costs at NNSA Laboratories	14
2.05 — A Common Workload Manager Lowers Barriers for Tri-Lab Scientists	16

Section 3: Research and Advanced Development

3.00 — Research Provides a Solid Foundation for Fulfilling the Laboratory's Mission	17
3.01 — New Modeling of NIF Targets Helps Protect Optics and Diagnostics	19
3.02 — High-Resolution Simulations Help Guide the Design of NIF Targets and Experiments	20
3.03 — New Maxwell Solver Enables Scalable Electromagnetic Simulations	22
3.04 — Proteus Offers Timely Analysis in the Intelligence Arena	23
3.05 — Climate Modeling Research and Enabling Technology Have Worldwide Impact	24
3.06 — Cyber Security Research Creates Innovative Tools to Protect Computer Networks	26

Section 4: Applications Software Development

4.00 — Computational Scientists Facilitate Programmatic Success	27
4.01 — NIF Integrated Computer Control System Software Demonstrates Full-Scale Operation	29
4.02 — NIF Automatic Alignment Shows Accuracy, Reliability, and Speed	30
4.03 — Automated Target Positioners Reduce the Setup Time for NIF Shots	31

4.04 — NIF's Laser Performance Operation Model Facilitates End-to-End Shot Planning and Analysis	32
4.05 — NIF Data Repository Manages Long-Term Storage and Provides Analysis and Visualization Tools	33
4.06 — Computation Provides Leadership to Handle Next-Generation Biothreats	34
4.07 — Argus Upgrades Safeguard National Security Facilities	36
4.08 — ALE3D Sets Its Sights on the Future	37
4.09 — A Context-Aware Nuclear Evaluation System Increases Security at Ports of Entry	38

Section 5: Institutional Support

5.00 — Computation's Institutional Contributions Make a Difference Across the Laboratory	39
5.01 — ASC V&V Best Practices Yield Better Software at Lower Cost	41
5.02 — Institutional Computing Advances Grand Challenge Science	42
5.03 — Academic Interactions Produce Research and Recruits	44
5.04 — Server Virtualization Reduces Costs in Weapons and Complex Integration Organization	45
5.05 — Cyber Security Protects Against Evolving Threats to LLNL's Computing Assets	46
5.06 — Production Begins on Institutional Wireless Network	47
5.07 — Process Automation Enables Centralized Desktop Configuration Management	48

Section 6: Appendices

6.01 — Computation Directorate Industrial Collaborators	50
6.02 — Academic Outreach	51
6.03 — Publications	56
6.04 — Acronyms and Abbreviations	61

1.00 — Extraordinary Computational Achievements Prevail during a Year of Transition

Dona Crawford, Associate Director, Computation Directorate



If there is a single word that both characterized 2007 and dominated the thoughts and actions of many Laboratory employees throughout the year, it is *transition*. Transition refers to the major shift that took place on October 1, when the University of California relinquished management responsibility for Lawrence Livermore National Laboratory (LLNL), and Lawrence Livermore National Security, LLC (LLNS), became the new Laboratory management contractor for the Department of Energy's (DOE's) National Nuclear Security Administration (NNSA). In the 55 years under the University of California, LLNL amassed an extraordinary record of significant accomplishments, clever inventions, and momentous contributions in the service of protecting the nation. This legacy provides the new organization with a built-in history, a tradition of excellence, and a solid set of core competencies from which to build the future.

I am proud to note that in the nearly seven years I have had the privilege of leading the Computation Directorate, our talented and dedicated staff has made far-reaching contributions to the legacy and tradition we passed on to LLNS. Our place among the world's leaders in high-performance computing, algorithmic research and development, applications, and information technology (IT) services and support is solid. I am especially gratified to report that through all the transition turmoil, and it

has been considerable, the Computation Directorate continues to produce remarkable achievements. Our most important asset—the talented, skilled, and creative people who work in Computation—has continued a long-standing Laboratory tradition of delivering cutting-edge science even in the face of adversity.

The scope of those achievements is breathtaking, and in 2007, our accomplishments span an amazing range of topics. From making an important contribution to a Nobel Prize-winning effort to creating tools that can detect malicious codes embedded in commercial software; from expanding BlueGene/L, the world's

most powerful computer, by 60% and using it to capture the most prestigious prize in the field of computing, to helping create an automated control system for the National Ignition Facility (NIF) that monitors and adjusts more than 60,000 control and diagnostic points; from creating a microarray probe that rapidly detects virulent high-threat organisms, natural or bioterrorist in origin, to replacing large numbers of physical computer servers with small numbers of virtual servers, reducing operating expense by 60%, the people in Computation have been at the center of weighty projects whose impacts are felt across the Laboratory and the DOE community.

The accomplishments I just mentioned, and another two dozen or so, make up the stories contained in this report. While they form an exceptionally diverse set of projects and topics, it is what they have in common that excites me. They share the characteristic of being central, often crucial, to the mission-driven business of the Laboratory. Computational science has become fundamental to nearly every aspect

LLNL systems appearing in the top 100 on the Top500 Supercomputers list (November 2007).

2007 Rank	2006 Rank	Linpack Speed	System Name	Description
1	1	478.2 TF	ASC BlueGene/L	106,496-node system built in partnership with IBM; used for capability computing.
11	4	75.7 TF	ASC Purple	1,532-node system built in partnership with IBM; used for capability computing.
29	New	36.6 TF	M&IC Atlas	1,152-node Linux cluster system used for unclassified capability computing.
38	New	27.3 TF	ASC Minos	864-node Linux cluster system used for classified capacity computing.
47	19	19.9 TF	M&IC Thunder	1,024-node Linux cluster system used for unclassified capability computing.
61	New	18.2 TF	ASC Rhea	576-node Linux cluster system used for classified capacity computing.

of the Laboratory's approach to science and even to the conduct of administration. It is difficult to consider how we would proceed without computing, which occurs at all scales, from handheld and desktop computing to the systems controlling the instruments and mechanisms in the laboratories to the massively parallel supercomputers. The reasons for the dramatic increase in the importance of computing are manifest. Practical, fiscal, or political realities make the traditional approach to science, the cycle of theoretical analysis leading to experimental testing, leading to adjustment of theory, and so on, impossible, impractical, or forbidden. How, for example, can we understand the intricate relationship between human activity and weather and climate? We cannot test our hypotheses by experiment, which would require controlled use of the entire earth over centuries. It is only through extremely intricate, detailed computational simulation that we can test our theories, and simulating weather and climate over the entire globe requires the most massive high-performance computers that exist. Such extreme problems are found in numerous laboratory missions, including astrophysics, weapons programs, materials science, and earth science.

Without doubt, the transition to LLNS has been challenging for both the organization and its people. Organizationally, the Laboratory has made


many changes that may appear to be simple or superficial but in fact represent wholesale changes in the way we operate. Largely, these have arisen as we moved from the public to the private sector and found that the laws governing the way we operate now are entirely different. We have had to reinvent our business methods from the ground up; having to conform to new laws meant developing entirely new systems for benefits, payroll burdens, work schedules, and exempt and non-exempt personnel

processes. The cost of this retooling is high, and at the same time, the Laboratory's budget was cut. These factors dictated cost reductions at all levels across the Laboratory. To this end, Computation has made significant reductions by consolidating office space and reducing energy consumption. For example, the energy used to cool the computer rooms has always been exempted from cost-cutting efforts, but Computation voluntarily reexamined its energy usage, raised the temperature in

the Terascale Simulation Facility machine rooms, and produced significant cost savings for the Laboratory, with no impact to our customers.

It is natural to ask, "Where do we go from here? What does the future hold for Computation?" These questions always dominate management, but in this time of transition, they have become the daily reality of all employees. The answer is simple—Computation will continue to provide the Laboratory programs with the machines,

The integrated computer control system for NIF monitors thousands of components, such as this calorimeter, which measures the energy of the laser.



algorithms, software, and services they need to meet their specific mission goals. Computation will provide unmatched IT support throughout the Laboratory. Perhaps most important, Computation will blaze a path into new terrain, extending the efforts of today into previously intractable areas and opening up unexplored fields of research.

The directorate's priorities are those of the Laboratory, and they are evident in our efforts to address LLNL Director George Miller's top priorities, delineated in his "A-List." Computation has a role to play in almost all of the items on the list; in several, however, we play a very central role. For example, the list calls

out continuing to apply Purple and BlueGene/L to "Grand Challenges" and securing a path forward for Sequoia, the next generation of supercomputer, which is designed to increase performance by an order of magnitude over existing machines. The A-List also specifies positioning the Laboratory to develop insights, technologies, and operational capabilities to counter proliferation, catastrophic terrorism, and asymmetric warfare. To this end, we are embarking on a major effort to understand data-intensive computing, because the scale of the datasets that must be handled in the intelligence domain is truly staggering. Entirely new methods, running on entirely new hardware, will be required to make headway, and Computation plans to be at the center of those efforts. The A-List instructs us to help NNSA develop a sustainable nuclear weapons enterprise, to help NIF meet the goal of attaining ignition in 2010, and to help develop solutions for international energy and environmental challenges. To do this, LLNL will rely on Computation to continue delivering top-notch hardware, algorithmic, and software support to the weapons codes, to perfect and maintain the massive control system that allows NIF to operate effectively, and to scale up climate models to the 100K processor size, permitting accurate modeling on a global scale. The Director's A-List also specifies that LLNL will ensure its competitiveness by significantly reducing the cost of doing business and improving the Laboratory's operational efficiency and effectiveness.


Computation's response is to propose centralized, rapid-response computer system administration. This would provide enhanced automation, leverage economies of scale, and create pools of deep expertise in particular areas. In short, Computation will be intricately involved in almost everything the Laboratory has declared to be important.

I have little doubt transition issues will continue to present challenges, as well as difficult and uncomfortable realities, during the next year. I am confident, though, that LLNL and Computation will continue harnessing the efforts of its extraordinarily talented workforce.

These are the realities that drive our efforts and generate an excitement among personnel about the potential for what we do. Computing helps us understand the world we live in and helps us to improve that world. We are at the center of momentous activity and can have a sizable effect on the future of the nation. As long as this remains true, the Computation Directorate will remain a thriving enterprise.



Dona Crawford
Associate Director, Computation



This quantum chromodynamics simulation, to calculate the thermodynamic properties of quarks (particles that make up protons and neutrons), illustrates how BlueGene/L is employed on fundamental grand-scale problems.



1.01 — Awards and Recognition

A Prize-Winning Organization

The stories in this annual report present a cross-section of Computation's work in high-performance computing, research, applications and software, and institutional support. In addition to the directorate-wide accomplishments those stories highlight, Computation projects and personnel received several significant awards in 2007, some of which are featured in this section.

Generally, the awards given to project teams are the culminations of years of strategic planning and intense effort, often involving more people than those who are included in the award citation. Similarly, the awards given to individuals tend to recognize sustained, sometimes career-length, records of extraordinary contribution.

The people who performed the prize-winning work share several characteristics and represent the very best that Computation has to offer: they are dedicated, persistent, innovative, and have a passion for the Laboratory's mission. They serve as role models for all Computation personnel and as ambassadors from Computation to the rest of the Laboratory and the broader scientific community.

A Nobel Effort in Climate Change

Nearly a dozen Computation scientists, software engineers, and system administrators were key contributors to the reports of the Intergovernmental Panel on Climate Change (IPCC) that won the 2007 Nobel Peace Prize with former U.S. Vice President Al Gore. The IPCC and former Vice President shared equally in

this prestigious honor for their efforts to increase and disseminate knowledge about human-induced climate change, and to lay the foundations for the measures that are needed to counteract such change.

Computation's involvement in the IPCC came by way of the Laboratory's Program for Climate Model Diagnosis and Intercomparison (PCMDI), which made major contributions to all of the IPCC reports, from the First Assessment Report in 1990 to the Fourth Assessment Report


in 2007. PCMDI was directly responsible for the success of the multimodel climate output database identified in the Fourth Assessment Report and cited in the Nobel Prize press release.

2007 Gordon Bell Prize

For the third year in a row, a team of scientists from LLNL and IBM earned the Gordon Bell Prize for pioneering simulations on BlueGene/L. Led by LLNL's Fred Streitz, the team included Kyle



Several Computation managers join in celebration with the directorate's contributors to the Nobel Prize-winning effort (from left): Velimir Mlaker, Robert Drach, Kyle Halliday, Dean Williams, Anthony Hoang, Associate Director Dona Crawford, Charles Doutriaux, Daniel Bergmann, Division Leader Peg Folta, Art Mirin, Department Head Patricia Damkroger, and Principal Deputy Associate Director Ronald Schmucker. Not pictured: Jennifer Aquilino, Robert Gibson, Renata McCoy, John Tannahill, and Michael Wickett.



Scientists can now predict how a wave forms, atom by atom. LLNL's Gordon Bell Prize-winning first-of-a-kind simulation of a Kelvin-Helmholtz instability in molten metals provides a much needed understanding of the motion of fluids and the interactions of the fluid with its boundaries.

Caspersen, Jim Glosli, David Richards, and Robert Rudd of LLNL, and John Gunnels of IBM. By performing extremely large-scale molecular dynamics simulations using an innovative computational technique, the team was able to study, for the first time, how a Kelvin-Helmholtz instability develops from atomic-scale fluctuations into micron-scale vortices. This work has important implications for such Laboratory research efforts as NIF's laser fusion experiments and developing applications for nanotube technology.



The *hypre* development team (from left): Van Emden Henson, Robert Falgout (Project Leader), Ulrike Meier Yang, Tzanio Kolev, Charles Tong, Allison Baker, Barry Lee, Jeffrey Painter, and Panayot Vassilevski.

BlueGene/L Continues Its Reign

The newly expanded IBM BlueGene/L system claimed the top spot on the Top500 list of the world's fastest supercomputers for the seventh time in a row, a distinction it shares with only one other machine—Sandia's ASCI Red. BlueGene/L now achieves a Linpack benchmark performance of 478.2 trillion floating-point operations per second (TF), compared to 280.6 TF six months ago, before its upgrade. (See Section 2.01.)

Teams of LLNL and IBM scientists also used BlueGene/L to win three of the four High-Performance Computing Challenges at the 2007 Supercomputing Conference, proving that the system is capable of solving a wide range of computational problems. BlueGene/L earned yet another honor in 2007 by ranking near the top of the Green500 list, a new compendium of the most energy-efficient supercomputers in the world.

Linear Solutions at *Hypre* Speed

Hypre, a software library of high-performance preconditioners and solvers developed by Computation scientists, received a 2007 R&D 100 Award from *R&D Magazine's* annual competition for the 100 most technologically significant products. *Hypre* is unique in its ability to provide solution algorithms that are effective on a spectrum of diverse problems, are easily accessible via multiple user interfaces, and effectively use the power of today's supercomputers to perform large simulations. Research institutions and private companies have used the free open-source software to simulate maxillofacial surgeries, magnetic fusion energy plasmas in tokamaks and stellarators, blood flow through the heart, fluid flow in steam generators for nuclear power plants, and pumping activity in oil reservoirs. (See Section 3.03.)

Lifetime Contributions and Leadership Excellence

Computer scientist **Nelson Max**, who recently accepted a full-time appointment at UC Davis, was awarded the Steven A. Coons Award by the Association for Computing Machinery's Special Interest Group on Graphics and Interactive Techniques (ACM SIGGRAPH). Max was recognized for his work in pioneering scientific visualization, his extensive technical contributions, and his generous participation in the encouragement and stimulation of ideas and intellectual exchange in computer graphics.

David Keyes, a computational mathematician and acting Director of LLNL's Institute for Scientific Computing Research, received the Institute of Electrical and Electronics Engineers (IEEE) 2007 Sidney Fernbach Memorial Award for his outstanding contributions to the development of scalable numerical

algorithms for solving nonlinear partial differential equations and for his exceptional leadership in high-performance computation. Keyes is world-renowned for contributions to Newton–Krylov–Schwarz methods for the efficient solution of nonlinear partial differential equations on high-performance computers. These methods are at the heart of many applications, including aerodynamics, radiation transport, acoustics, and magnetohydrodynamics.

Computer scientist **Dean Williams** received *Science Spectrum* magazine's Emerald Award for his technical leadership and contributions to the international climate community. Williams leads a variety of software projects aimed at data analysis and visualization of climate models. He is also active in LLNL's recruiting efforts and participates in the summer student program to help ensure that the Laboratory continues to hire qualified minorities into scientific fields.



Nelson Max



David Keyes



Dean Williams

2

Computation Develops Strategies for Preeminent National Security Computing

For practitioners of science in the national security enterprise, the future is a complex and dimly lit landscape, fraught with challenges and opportunities, known and unknown. For Computation, and particularly the Livermore Computing (LC) Center, to continue to play a pivotal role as in years past, it is essential to follow sound and adaptive centralized computing plans. Five interrelated strategies form the backbone of Computation's approach to increasing productivity while ensuring continued scientific distinction.

First, it is essential to maintain a balance across platform investments. Computational studies depend on a spectrum of calculations, from large numbers of smaller calculations (capacity computing) to smaller numbers of large or even heroic calculations (capability computing). The recent BlueGene/L, Peloton, and Tri-Laboratory Capacity Cluster (TLCC) procurements reflect this reality; by mid-2008, LC will field nearly 730 trillion floating-point operations per second (TF) in capability resources (Purple, Atlas, and BlueGene/L computers) and 370 TF in capacity systems. Since capability resources are shared with other laboratories,

2.00 — Computation Develops Strategies for Preeminent National Security Computing

CONTACT INFORMATION Michel McCoy, mgmccoy@llnl.gov



the ratio of capacity to capability is approaching a very viable 1:1 for Lawrence Livermore National Laboratory (LLNL).

Second, LLNL must maintain its technical leadership in high-performance computing (HPC). Today, Department of Energy Headquarters evaluates its contractors on meeting mission requirements and on business efficiency. In 2007, LC tallied several remarkable achievements on both fronts. A 60% expansion and upgrade brought BlueGene/L to a peak 590 TF, and, just seven months after the initial expansion funding request, LLNL scientists used the expanded BlueGene/L to win a fourth Gordon Bell Prize. LC operated Purple as a national user facility, and, throughout 2007, the machine enjoyed a lofty 94% utilization rate, with but 5% idle time and 1% down time. LLNL was selected to coordinate the TLCC procurement and to provide the future Tripod Operating System Software (TOSS). The extraordinarily low \$62K-per-TF procurement cost and savings from common use of TOSS will ensure that TLCC is a low-cost, reliable solution for the National Nuclear Security Administration (NNSA). LC developed a cost-shared (with the institution) plan for upgrading LLNL's Terascale Simulation Facility to provide 30 megawatts (MW) of power for computers in that building by 2011. The first milestone was completed, bringing the building from 12 MW to 15 MW. Working with the Advanced Simulation and Computing (ASC) Program on the \$250M Sequoia procurement, LC has submitted Critical Decision 1 to Headquarters. The plan would see an "initial delivery system" in late 2008 and the final system in 2011.

The third strategy focuses on providing first-class services to the scientists. LC strives to balance an operations and services environment. LC maintains a highly effective hotline and provides specialized services to customers across the programs at LLNL, the Tri-Lab community (Lawrence Livermore, Los Alamos, and Sandia national laboratories), and the ASC Alliances (LC historically provides more than 80% of the cycles to the Alliances). These services include applying advanced tools for performance, memory management, and visualization, and giving dedicated application times (clearing the machines for

many-hour full-system calculations at night). Occasionally, subtle bugs render important codes unusable; LC dedicates personnel to work the most difficult issues—sometimes for many weeks—until resolution is achieved.

Fourth, LC is dedicated to providing high-quality services in both the classified and unclassified realms. While the bulk of LLNL's resources are dedicated to classified work, LC nevertheless provides a world-class unclassified computing environment. This effort is supported by LLNL and ASC, who recognize that maintaining exceptional quality of the science and technology base at LLNL necessitates fostering open publication of scientific results and collaboration with academia. LLNL's Multiprogrammatic and Institutional Computing Program supports unclassified computing services, as demonstrated by its procurement of a 40-TF replacement for the Thunder computer and its shared investment in the upgrade of BlueGene/L, from which 230 TF will become unclassified in 2008.

To effectively use this robust unclassified environment, LC managers are developing the fifth strategy: working to enlarge LLNL's portfolio of work for others (WFO). Initial steps include making a co-investment with a group of major vendors, which will facilitate WFO capability by creating an open 100-TF scalability test platform. This machine, to be shared with the co-investors, is for developing system software (file systems, operating systems, tools, etc.) at massive scale. Adoption of large-scale scientific and engineering codes by industry has been impeded by "chicken-and-egg" issues regarding availability of scalable codes and platforms. Using the Peloton model with an open software stack and having an open platform for hardening system software will enable LLNL and industry to build truly scalable scientific applications together. This capability will allow LLNL to have an impact on American technological competitiveness and will open future funding possibilities in computational science to our scientists and engineers.

The stories that follow highlight how LC applied the first four strategies in 2007 and hint at what may be expected from all five strategies in the future. This tale is not one of LC working alone; the efforts of LC are visible in stories throughout this report and in work conducted across the Laboratory.

2.01 — Expanded BlueGene/L Brings Predictive, Fully Integrated Calculations Closer to Reality

CONTACT INFORMATION Kim Cupps, cupps2@llnl.gov



OVERVIEW

With a peak performance of 367 TF, the original BlueGene/L system, which LLNL accepted in December 2005, was intended primarily for materials science simulations. Between May and September 2007, BlueGene/L was expanded by roughly 60%, enabling users to address a broader range of weapons science calculations. In November, BlueGene/L appeared a record-tying seventh time at number one on the Top500 list of the world's fastest computers. At 596-TF peak performance, it is approximately 2.9 times faster than the next fastest computer.

Increasingly important, BlueGene/L is quite "green," consuming less than one-third of the power per TF used by a non-BlueGene supercomputer. Although the

system expansion was a high-risk endeavor with much uncertainty, the final result is positive and impressive.

PROGRESS IN 2007

When LLNL asked IBM if it was technically possible to use the company's new high-memory nodes while expanding BlueGene/L by 60%, IBM engineers noted some concerns. The hardware was designed to scale to greater than the 64 original 1,024-node racks. However, the original software for the control system had been designed and tested for a maximum of 65,536 nodes (the original size of BlueGene/L), not the proposed 106,496 nodes. IBM engineers predicted that the control software could be scaled given a

sufficient development effort, but significant limitations in the software interface were likely. Any solution delivered in the proposed five-month time frame would be a bare-bones upgrade, at best.

LLNL decided to proceed with the expansion and created an integration and delivery schedule. IBM began modifying software and developed a manufacturing pipeline for the new racks. The computer room in the Terascale Simulation Facility was modified during May and June to make space for 40 contiguous new racks of equipment. Modifications, such as adding floor supports and cable trays and increasing the power distribution for the new racks, were completed with no significant impact on use of the supercomputer.


Rack installation began in June. Then in late summer, the IBM-LLNL team encountered several difficulties, including

damaged racks and problems with certain third-party-supplied components. As a result, the team had to remanufacture several hundred cards and three complete racks. Fortunately, these difficulties caused only a minor delay in the delivery schedule.

To bring up the new racks, Livermore's system administrators ran numerous tests on a partially complete 98,304-node system, often working around the clock in the weeks-long process. They identified a small number of software defects, which IBM quickly resolved. The final eight racks arrived in September, and after another month of intensive integration and software testing, the IBM-LLNL team made the expanded BlueGene/L system available to a set of early users. In early October, teams from Livermore and Los Alamos, working alternately, ran simulations using an unprecedented 212,992 processors on 106,496 nodes.

SIGNIFICANCE

Despite IBM's conservative prediction, software for the expanded BlueGene/L is fully functional and scalable, offering users a sufficient set of features. With BlueGene/L's improved capabilities, computational scientists can now scale weapons science calculations to record processor counts, improve scaling of integrated weapons codes, and dramatically enhance the performance of complex models and algorithms in macroscopic-scale simulation codes. Key scientific studies on the expanded system have, for example, established the connection between atomistic behavior and practical issues in stockpile stewardship, allowing LLNL scientists to move ever closer to the goal of running predictive, fully integrated weapons calculations in three dimensions.



The expanded BlueGene/L system has a peak performance of 596 TF, about 2.9 times faster than the second machine on the Top500 list of high-performance supercomputers.



2.02 — BlueGene/L Enables Large-Scale Science Runs

OVERVIEW

Delivering on LLNL's critical mission responsibilities and discovering the underlying fundamental science drive the Laboratory's quest for increasingly higher resolution in simulation. The IBM-built BlueGene/L computer serves as LLNL's largest resource for addressing issues connected with extreme resolution. As originally procured, BlueGene/L consisted of 131,072 processors (on 65,536 two-processor nodes) and a total of 33.7 terabytes of memory.

BlueGene/L was initially used to address urgent questions associated with plutonium aging. Its simulations provided quantitative information characterizing the effects of self-irradiation damage and provided quantitative bounds on the effects of such damage for the U.S. stockpile. Originally, BlueGene/L was focused on science codes, which typically address a single physics process. However, early efforts with integrated design codes, which couple many types of physics processes together, also showed good parallel scaling efficiencies on BlueGene/L, thereby expanding the role BlueGene/L plays at LLNL.

Atomistic development of instabilities in two dimensions.


PROGRESS IN 2007

BlueGene/L's 2007 schedule was fully subscribed and included three very large efforts: the Los Alamos W76 Life Extension Program, the National Ignition Facility phase-plate final design, and an ASC material strength Level 2 milestone. During these calculations, the mean time between failures was two weeks, which

is a major achievement for a 360-TF computer.

BlueGene/L's success led to a decision in 2007 to use a combination of ASC Program and LLNL institutional funding to increase the size of BlueGene/L to 212,992 processors and apply it to foundational science discovery, while simultaneously preparing codes to

operate at the soon-to-be-available scale of one quadrillion floating-point operations per second (petaFLOPS). Among the first full-scale problems on the expanded BlueGene/L were two classical molecular dynamics simulations. One simulation was a Kelvin–Helmholtz instability at the interface between molten copper and






CONTACT INFORMATION Dale Nielsen, dnielsen@llnl.gov

aluminum in two dimensions. Next year, a three-dimensional calculation will more completely reveal the nature of the interface dynamics, connecting the finest-scale hydrodynamic simulations and the atomic scale.

In 2008, BlueGene/L will be divided, and 81,920 processors will be moved to the unclassified network, greatly

increasing the number of researchers at LLNL who can enlarge the size of their simulations. Of the processors added this year that will remain on the classified system, 65,536 contain double the normal memory, which expands the scope and complexity of the integrated design code problems that can run on BlueGene/L.



Shocks through copper into vacuum (top) and helium (bottom).

SIGNIFICANCE

The importance of high-resolution simulations cannot be overstated, because such studies scientifically establish the connection between atomistic behavior and practical issues in stockpile stewardship and other Laboratory missions. Many other Tri-Lab, LLNL programmatic, and basic science simulations are scheduled for the first half of 2008, while BlueGene/L is still in its fully expanded size. The LLNL vision of the BlueGene architecture as an affordable, low-power path to extreme performance will continue to broaden its impact in 2008 and beyond.



2.03 — Next-Generation Development Environment Tools Pave the Way to Sequoia-Scale Computing

OVERVIEW


The Integrated Computing and Communications Department’s (ICCD’s) Advanced Systems Architect Mark Seager posed a provocative question to ASC’s Application Development Environment and Performance Team (ADEPT) in 2007. “How,” he asked, “are you going to debug an application running on one million processors?” His question is not far-fetched technically; it is a challenge that is relevant today. Current plans for the Sequoia system

architecture call for more than one million central processing unit cores, and the target deployment date is 2011—just four years away. The ADEPT team, which includes personnel from ICCD’s Development Environment Group and from the Computing Applications and Research Department’s Center for Applied Scientific Computing, responded to Seager’s challenge by pursuing a multifaceted strategy to provide an ultra-scalable tool suite for correctness and performance debugging.

While debugging at the Sequoia scale is ADEPT’s grand challenge, it is only one of many current and future requirements. The team’s guiding philosophy for code development on HPC systems is embodied in the “Livermore model,” in which the Message Passing Interface (MPI) library is the dominant parallel programming paradigm. But a growing need for parallelism based at the node level led code developers to incorporate more OpenMP and other threading methods into the programming

model. The ADEPT team responded with tools that address these changes. The team used open-source collaborations and formed strategic partnerships with commercial vendors to provide best-in-class, consistent tools that can be used on all platforms, from desktops to machines at the Sequoia scale. Responding to input from LLNL code teams, the team focused on creating improved memory-usage-tracking tools and memory-correctness debuggers, among other products.

The predicted performance of the Stack Trace Analysis Tool (STAT) at BlueGene/L scales, with an estimated latency of 2.6 seconds and 131,072 tasks, in Co-Processor (CO) mode and Virtual Node (VN) mode using various tool communication tree depths. As application task-counts increase, additional layers can be added to STAT’s communication tree to achieve continued scalability.



CONTACT INFORMATION

W. Scott Futral, futral2@llnl.gov
Bronis de Supinski, bronis@llnl.gov



PROGRESS IN 2007

ADEPT partnered with TotalView Technologies to improve the scalability of the TotalView debugger for current and future LLNL computer systems. To address the Laboratory's immediate needs for Linux clusters based on the InfiniBand interconnection switch, the team improved the method by which target applications handle dynamic library loading. They tested the method on an application code running 3,072 tasks on the Atlas cluster, and the start-up time improved from 160 to 2.5 seconds. The team also investigated viable design options that will advance the TotalView debugger scalability to the level demanded by the Sequoia system.

Understanding how memory is used in large-scale parallel applications is essential for debugging and for efficient system use. To address this need, ADEPT is developing the lightweight and flexible memP heap-profiling library. Application developers will use memP when conducting large application runs to identify the tasks that consume the greatest amount of heap memory and the location in source code of the allocation calls. Although it is still early in the development process, memP is based on

the highly successful mpiP profiling tool, which has run tens of thousands of tasks on the BlueGene/L system. The memP profiler is currently available on LLNL's x86_64-Linux clusters and will soon be ported to AIX operating systems.

ADEPT also initiated and supervised three memory tools projects that were funded or co-funded by LLNL in 2007. These projects included deploying a port of the open-source Valgrind Memcheck tool to BlueGene/L, developing a highly scalable batch mode for TotalView's lightweight memory-debugging tool, and revising the Valgrind-based thread-checking tool Helgrind to address an expected growth in the use of threading in LLNL's large-scale HPC applications.

Among the ADEPT team's activities is their participation in Open|SpeedShop, a collaboration on open-source tools that partners the Tri-Labs with the Krell Institute. In 2007, members of the collaboration incorporated improvements into Open|SpeedShop that were identified in comprehensive usability studies conducted by code teams from Lawrence Livermore and Los Alamos national laboratories. Improvements were made to MPI tracing capabilities, parallel job control, and postmortem

analysis of new performance experiments. ADEPT and their Open|SpeedShop partners are developing a new hierarchical communication infrastructure to provide efficient performance analysis, even at Sequoia scales.

ADEPT also investigated highly innovative approaches for novel tools and tool infrastructures aimed at future ASC platforms. The P^NMPI infrastructure allows the concurrent execution of multiple MPI tools and supports their interoperability. It enables modular tool development, simple and quick application-specific tool prototyping, and the virtualization and multiplexing of existing tools such as profilers or message checkers. Overall, P^NMPI provides unprecedented flexibility in the design

and deployment of MPI tools, allowing developers to target specific performance problems more quickly and efficiently.

An important part of ADEPT's Sequoia-debugging strategy is to use lightweight tools to reduce the problem search space to a scale that can be easily managed by full-featured tools. For example, the Stack Trace Analysis Tool (STAT) merges application stack traces to identify process equivalence classes. A representative task from each equivalence class can be fed into a full-featured debugger for root-cause analysis. STAT has already demonstrated the merging of latencies less than one tenth of a second for 4,096 tasks on several LLNL Linux clusters. ADEPT also created a new tool called STATBench to emulate the performance of STAT on arbitrary machine architectures.

SIGNIFICANCE

Advances in the capabilities of system and debugging tools for current LLNL systems are already paying dividends in terms of greater user productivity and efficiency. For example, the Valgrind Memcheck tool now works in the presence of MPI programs and has identified elusive memory-overwrite problems in a production ASC code, which had eluded programmers for several months. By advancing tools infrastructure and scalability with STAT and P^NMPI, ADEPT will help meet the challenge of providing Sequoia-scale tools and debugging capability.



2.04 — Common Capacity Computing Lowers Costs at NNSA Laboratories

OVERVIEW

In the past few years, NNSA's Tri-Lab community has made a major effort to standardize capacity-computing hardware and software. Capacity-computing systems provide efficient, cost-effective computing power to solve many small- to mid-sized problems or to prepare for runs on a capability system. In contrast, capability computing (such as the work done on BlueGene/L) uses the maximum power of a supercomputer to solve a very large problem in the shortest amount of time. The problems being run on a capability system are often so large or complex that no other computer can process them.

Although capability computing garners the plaudits and honors, capacity computing is a vital component of LC's computational environment. Capacity computing represents nearly 30% of the available resources on LC's classified network and more than 50% on the unclassified network. Even so, LC users regularly complain that insufficient capacity-computing cycles are available.

To meet the growing needs for cost-effective capacity computing, the ASC Program developed the TLCC procurement strategy. The TLCC strategy is to build a common hardware environment across the three laboratories,

capitalizing on inexpensive commodity Linux clusters. LLNL, which has fielded several large-scale, high-performance Linux clusters, leads the NNSA effort to develop TLCC requests for proposals.

The ASC Program had previously chartered Tripod, a common software strategy for Linux capacity clusters. Tripod's goal is to develop a seamless software environment for use by the tri-laboratory community. Tripod working groups are defining compilers and application development tools, system software, reliability and serviceability metrics, the user environment, and computer-monitoring capabilities.


With TLCC, Tripod's software strategy can now be based on a truly common hardware platform.

PROGRESS IN 2007

After a competitive bidding process, NNSA awarded the TLCC procurement contract to Appro in September 2007. With a total value of \$26M, the contract yields an exceptionally low cost of approximately \$62K per TF. The general architecture of the procurement strategy is modeled after LLNL's 2006 Peloton procurement. Peloton acquired the Zeus and Atlas machines for unclassified computing under the LLNL Multiprogrammatic and Institutional Computing Program, and Rhea and Minos for classified computing under the aegis of the ASC Program. The TLCC architecture is based on a scalable unit of 144 nodes with an InfiniBand interconnection switch. Each unit has a peak computing speed of 20.3 TF. System deliveries will begin in early 2008.

In 2007, the ASC Program issued two solicitations to the three laboratories. The first sought project proposals in broad areas to help develop a common capacity computing environment. Four of the 19 proposals submitted were selected for funding, and all four involve tri-laboratory collaboration. Los Alamos will lead projects on application monitoring, the shared-project work environment, and the Gazebo test and analysis suite. Livermore will lead the Open|Speedshop collaboration.

The Minos Linux capacity cluster, which is part of the Peloton procurement, has 864 nodes, 6,912 central processing units, and an operating speed of 33 TF. Shown here (from left) are Livermore computer scientists Trent D'Hooze and Kim Cupps.



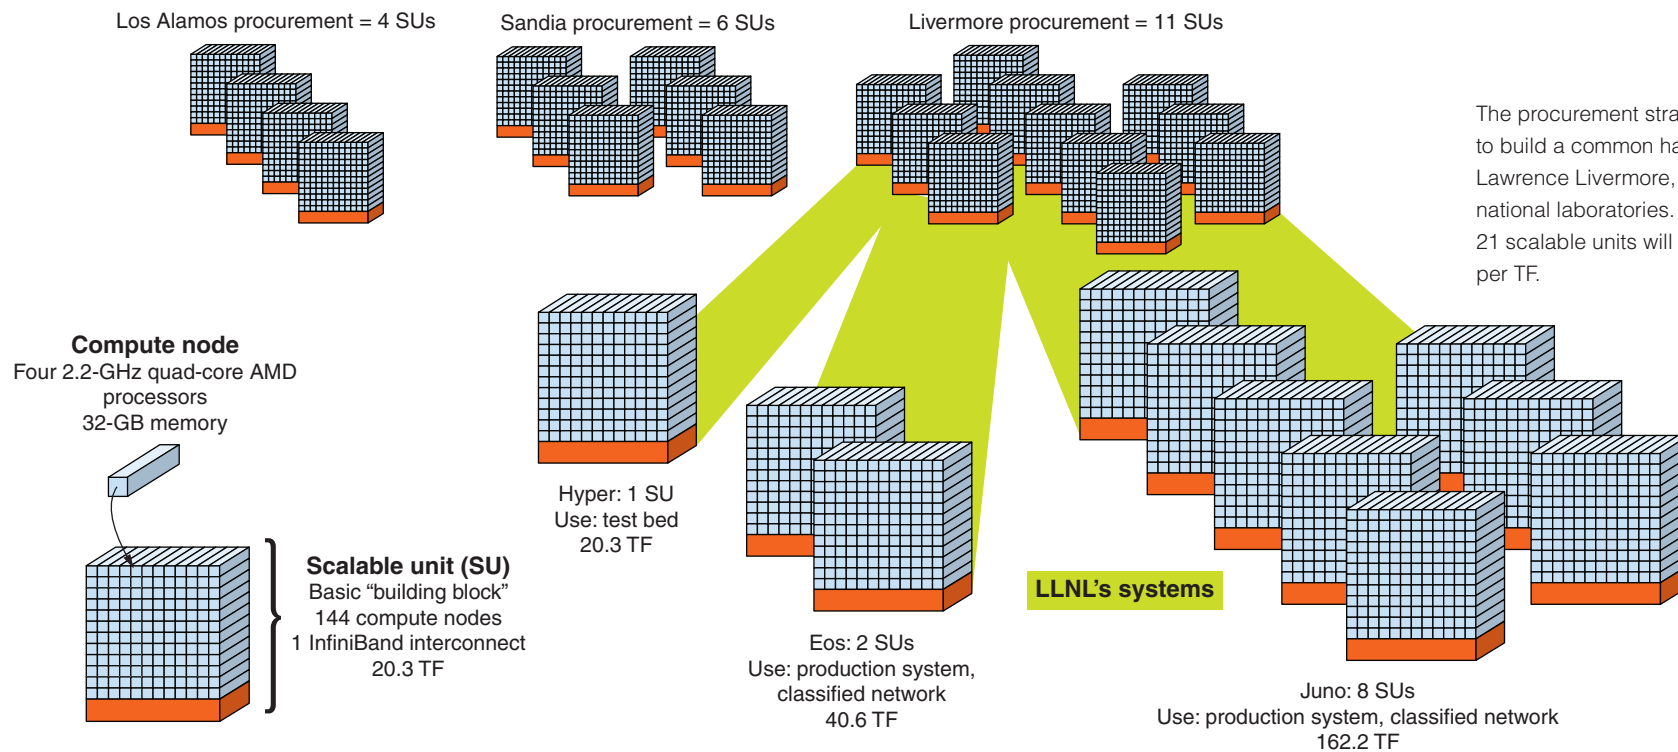


The second call solicited proposals for a common operating system software stack. Livermore teamed with Sandia to submit the winning proposal for TOSS, which will be based on the Clustered High-Availability Operating System (CHAOS) distribution running on LC's production Linux clusters. In addition to defining a Linux distribution, the proposal also specifies processes for packaging, delivery, ongoing support, and future evolution of the distribution.

Subsequent to the award, Livermore scientists worked with their colleagues at Los Alamos and Sandia as well as the ASC program managers to define a TOSS governance model. To the extent possible, the processes defined in the governance model mirror those developed to manage the CHAOS distribution at Livermore. The major challenge to the long-term success of TOSS will be ensuring that the three laboratories are equally involved in the support and evolution of the software's distribution.

SIGNIFICANCE

In the past, each laboratory procured its own capacity computing resources and ran a laboratory-specific operating system and user tools. The different hardware and software systems often made the computing environment challenging to users, which increased the difficulty of cross-laboratory collaborations and increased the overall costs to the NNSA complex. TLCC and TOSS offer the potential for significant cost savings not only in acquiring and supporting capacity systems, but also in reducing the effort and costs associated with porting applications to the different supercomputers at the three laboratories.



2.05 — A Common Workload Manager Lowers Barriers for Tri-Lab Scientists

CONTACT INFORMATION Don Lipari, lipari1@llnl.gov



OVERVIEW

Since 1990, the HPC systems within LC have used a proprietary workload manager called Livermore Computing Resource Management System (LCRM), which was written by members of the Computation Directorate's Integrated Computational Resource Management Group (ICRMG). The other two laboratories in the ASC Tri-Lab alliance (Sandia and Los Alamos) relied on commercially produced workload managers. Because the workload manager is the user's primary interface for scheduling

jobs across HPC clusters, having more than one workload manager running on the systems was confusing to users that ran jobs at multiple sites. In addition, each workload manager reported different machine usage and utilization statistics to ASC Headquarters.

In 2005, Tri-Lab representatives began working to establish a common workload manager to solve the inconsistencies, and, in 2006, they decided to purchase the Moab Workload Manager from Cluster Resources, Inc. (CRI).

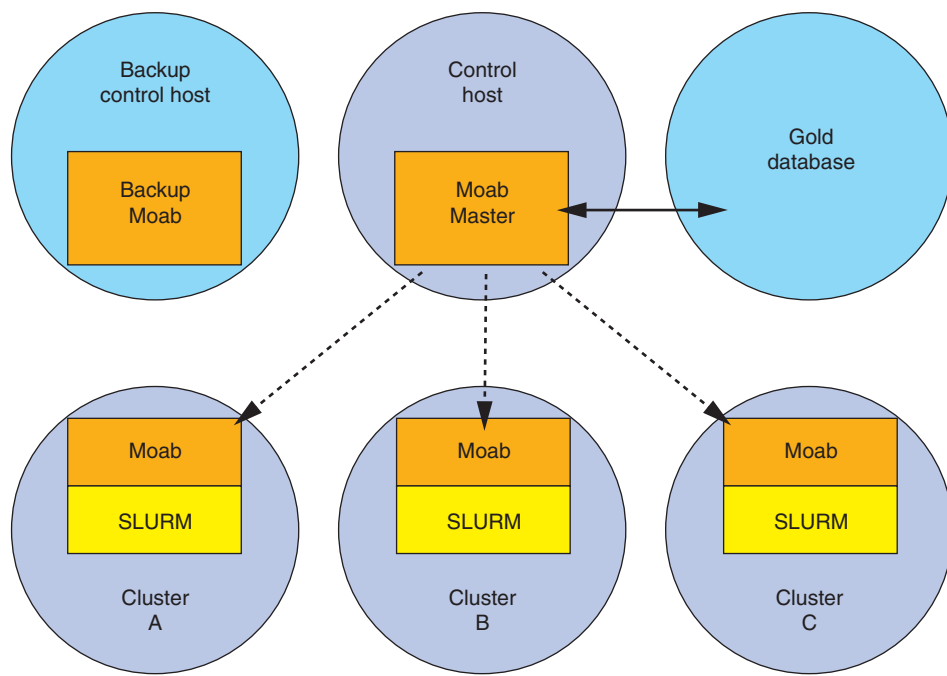
PROGRESS IN 2007

The decision to use Moab as a common workload manager at all three laboratories presented an opportunity for collaboration but also raised new challenges. To help chart a new course, the three laboratories and CRI adopted a support agreement, which defined the terms of support and the procedures for ensuring that requests from the three laboratories could be fairly prioritized and used to guide Moab product development.

As part of the agreement, ICRMG specified what enhancements needed to be made for Moab to support the batch-scheduling model on which LC users depended. Considerable improvements to Moab were also needed to interface to the Simple Linux Utility for Resource Management (SLURM). (SLURM is an open-source cluster-specific utility, originally developed by ICRMG, for launching jobs across clusters of all sizes, including BlueGene/L.) LLNL experts worked closely with CRI to tighten Moab's security to meet the standards imposed by LC's secure network. ICRMG

also wrote a set of test scripts that they used to validate each new release that was downloaded from CRI, as CRI responded to the issues and delivered an increasingly improved product.

As the various improvements were made, ICRMG and CRI created tools to provide a veneer around the Moab environment that emulated the previous workload manager, LCRM. This way, even though the new workload manager was managing their jobs, users could continue to submit and check the status of their jobs as they always had. The LCRM emulation was so successful that some users were unaware that the underlying workload manager had changed. ICRMG then wrote a set of tutorials to help users transition at their own pace to the native Moab commands. ICRMG and training staff from Computation's Services and Development Division developed training seminars for system administrators and users across LLNL. Throughout the transition, Moab administrators from the three laboratories worked together to develop uniform ways of configuring and running Moab.




Grid model showing how Moab ties a number of clusters together.

SIGNIFICANCE

The adoption of a common workload manager, chartered by ASC Headquarters, is the first step in building a Tri-Lab collaboration for delivering a common computing environment. LLNL met its ASC Level 2 milestone of running Moab on the first production cluster midway through 2007. In 2008, ICRMG will extend Moab to control BlueGene/L.

return
to TOC



3

Research Provides a Solid Foundation for Fulfilling the Laboratory's Mission

Innovations in computer science, computational algorithms, and information technology continue to serve as a foundational element for enabling science and engineering at Lawrence Livermore National Laboratory (LLNL). Recent advances in hardware have led to machines with hundreds of thousands of processors capable of handling terabytes of data. Coupled with the hardware growth, crucial and startling advances in the capability and scientific fidelity of multiphysics codes has led, for many applications, to the development of computational simulation as a true peer to theory and experiment. Frequently, the research required to create these innovations takes many years to mature, requiring vision and perseverance to push the technologies to fruition and implementation. This has been true, for example, of Computation's work in novel computational mathematics, numerical methods, and associated software libraries, which are being adopted in several programs at LLNL.

At the other end of the spectrum, rapid changes in technology and fast-evolving political or technological threats generate a class of research characterized more as "riding the advancing wave" than the traditional slow gestation described above. This category of work requires a different approach,

3.00 — Research Provides a Solid Foundation for Fulfilling the Laboratory's Mission

CONTACT INFORMATION John Grosh, grosh1@llnl.gov



usually involving rapid adaptation and the ability to integrate research into application components at a systems level. Results are often produced under the gun of necessity—implementations are needed in the field to meet rapidly evolving conditions. It is this dichotomy of approaches that makes research in Computation challenging and interesting.

Working in either research mode, Computation scientists constantly seek to influence and facilitate how the Laboratory meets its mission-specific goals. For example, the challenges facing the National Ignition Facility (NIF) represent a key opportunity for simulation to influence theory and experiment in fusion research. LLNL scientists are developing multiscale modeling for high-velocity debris and shrapnel effects due to the interaction of the lasers with the NIF target, to minimize the damage to optics and diagnostics equipment. This modeling capability was enabled by long-term research that integrated adaptive mesh refinement (AMR) with the more standard LLNL methods based on Arbitrary Lagrangian–Eulerian (ALE) hydrodynamics. This is the first integration of AMR and ALE into simulation code and has been an ongoing research project within the Center for Applied Scientific Computing (CASC) for the last six years. In turn, the ALE–AMR research builds on work in AMR techniques that started more than eight years ago when CASC researchers developed the scalable Structured AMR Application Infrastructure library.

Other examples of computation research impacting mission goals abound. LLNL has a longstanding commitment to inertial confinement fusion (ICF). Computation participates in this effort by contributing to the development of the simulation code HYDRA, which is used to design and analyze a variety of targets for NIF and other ICF applications. During the past decade, Computation scientists have worked closely with their counterparts in the Weapons and Complex Integration Principal Directorate to develop this scalable multiphysics code to simulate the full ignition dynamics in three dimensions. In recent years, simulations in support of NIF have consumed a major portion of the computational resources in Livermore Computing. Thus, research in highly efficient and scalable algorithms in support of the HYDRA code has been instrumental in enabling the Laboratory to perform such simulations to meet programmatic deadlines for NIF.

Two areas where Computation has a particularly strong international footprint are linear solvers and climate modeling and analysis. Recent research, built on long-term foundational work, has extended and enhanced the linear solver library *hypr*, enabling *hypr* to win a 2007 R&D 100 Award. *Hypr* is also being used in applications that focus on solving the Maxwell equations used in problems of electromagnetics simulation. In climate modeling, Computation scientists work with external research organizations, such as the National Center for Atmospheric Research, to model capabilities and substantially improve the resolution of the Community Climate System Model. In both research activities, Computation scientists continue to push the envelope of performance by exploring and developing scalable algorithms for systems such as BlueGene/L.

In support of the Office of Naval Intelligence, Computation and the Global Security Principal Directorate jointly developed a prototype information system to analyze ship traffic and open-source data. Great scientific agility was required to develop new algorithms for entity extraction. Teaming with industry partners was essential to adapt novel computer systems and solutions, and close collaboration with users (still ongoing) was necessary to build an effective end-to-end analysis system. This work provides valuable experience in building analytic systems and motivating future Computation research into computer and storage architectures and new algorithms in support of national security data-intensive applications.

Computation has begun to emphasize the computer science research required to advance cyber security. A huge challenge in this area is determining whether large software packages contain malicious code. Building on the ROSE compiler tool developed and used for source code analysis, Computation scientists are creating a framework to analyze large source-code and binaries. This novel work has enabled LLNL to gain access into the basic research community that supports the Department of Homeland Security (DHS) and the Department of Defense (DoD) work in information assurance and cyber security.

The wide range of research conducted in the Computation Directorate is represented in the following selection of stories; in the larger sense, the stories serve to illustrate the evolving nature of the challenges facing those who use science to help protect the nation.

3.01 — New Modeling of NIF Targets Helps Protect Optics and Diagnostics

CONTACT INFORMATION Alice Koniges, koniges1@llnl.gov



OVERVIEW

The Computation Directorate developed and deployed a novel, multiscale simulation code that has become an integral part of the NIF target design and evaluation process. Each NIF experiment will expend the target, so the target design must be carefully analyzed to ensure that shrapnel from the target will cause minimal damage to optics and diagnostics. The solution was to devise an innovative simulation regime with predictive capability for determining the debris and shrapnel effects. A team

of computer scientists developed a new three-dimensional (3D) parallel code that combines AMR with more conventional methods based on ALE hydrodynamics to perform advanced modeling of each target design. AMR provides true multiscale simulation that employs specialized physical models on different and appropriate scales, resulting in a hierarchical material model capability.


PROGRESS IN 2007

A high-powered laser chamber contains a significant amount of open space in a vacuum environment that surrounds the small fusion target—often a hohlraum—located at the chamber center. One of the team's goals is for the simulation code to model how the target disintegrates after being hit by either laser beams or by x-rays that result from the lasers interacting with other target components. For the purposes of the simulation, the target includes not only the hohlraum where the laser is


focused, but also ancillary target elements such as cooling rings, shields, or appendages that improve diagnostic capabilities. Pieces of the target that are nearest to where the laser is focused will be vaporized and, thus, are relatively benign. However, target components that are farther from the main laser focus point are subject to lower levels of energy and, therefore, may be fragmented. It is important to determine the size of these fragmented pieces and their velocity vectors after the laser shot so that optics and diagnostics that line the chamber will be protected from damage. The new code is unique in its ability to combine ALE with AMR to model the open space, as well as its ability to provide a choice of material damage models. The code includes damage and failure models, improved interface reconstruction, and input frameworks for advanced physics models. In addition to the development of new algorithms, a method to input target geometries in the ALE-AMR grid was formulated.

SIGNIFICANCE

The NIF ALE-AMR code is the most important design tool for determining if a target will have appropriate debris and shrapnel production. It is also used in design decisions for materials and geometries of target components in NIF and other high-powered laser systems.



(a and b) A recent calculation of two different cooling ring designs on an OMEGA laser target. A ring of metal is impulsively loaded from the center to demonstrate the fragmentation pattern of target cooling rings. (a) The ring is symmetric and loaded from the interior. (b) An alternative design has a notched ring. The inset shows the formation of potentially damaging pieces during the fragmentation process.


[return to TOC](#)

3.02 — High-Resolution Simulations Help Guide the Design of NIF Targets and Experiments

OVERVIEW

The radiation hydrodynamics code HYDRA is routinely used to design and simulate a wide variety of ICF targets, including ignition targets for NIF. When NIF's 192 laser beams are operational in 2010, they will deliver up to 1.8 megajoules of ultraviolet laser energy and 500 terawatts of power to the target. The baseline target design for ICF experiments on NIF uses a small metal cylinder, called a hohlraum, which surrounds a spherical plastic or beryllium fusion capsule containing deuterium–tritium fuel. The laser beams deposit much of their energy on the hohlraum wall, generating

intense thermal x-rays that illuminate the target capsule. These x-rays ablate the surface of the capsule, creating a rocketlike effect, which causes the capsule to implode. The implosion produces temperatures and pressures high enough to ignite the fuel by self-sustained nuclear fusion.

HYDRA can simulate the full ignition dynamics in three dimensions, including the hohlraum, capsule, and all relevant features. It resolves hydrodynamic instabilities that occur as the capsule implodes and intrinsic drive asymmetries that result from the laser illumination pattern. The flexible laser module can accurately represent illumination asymmetries resulting from

errors in laser beam power balance and aim. HYDRA also models radiation, electron, ion, and charged-particle transport, all of which are fully integrated with the laser drive and hydrodynamics.

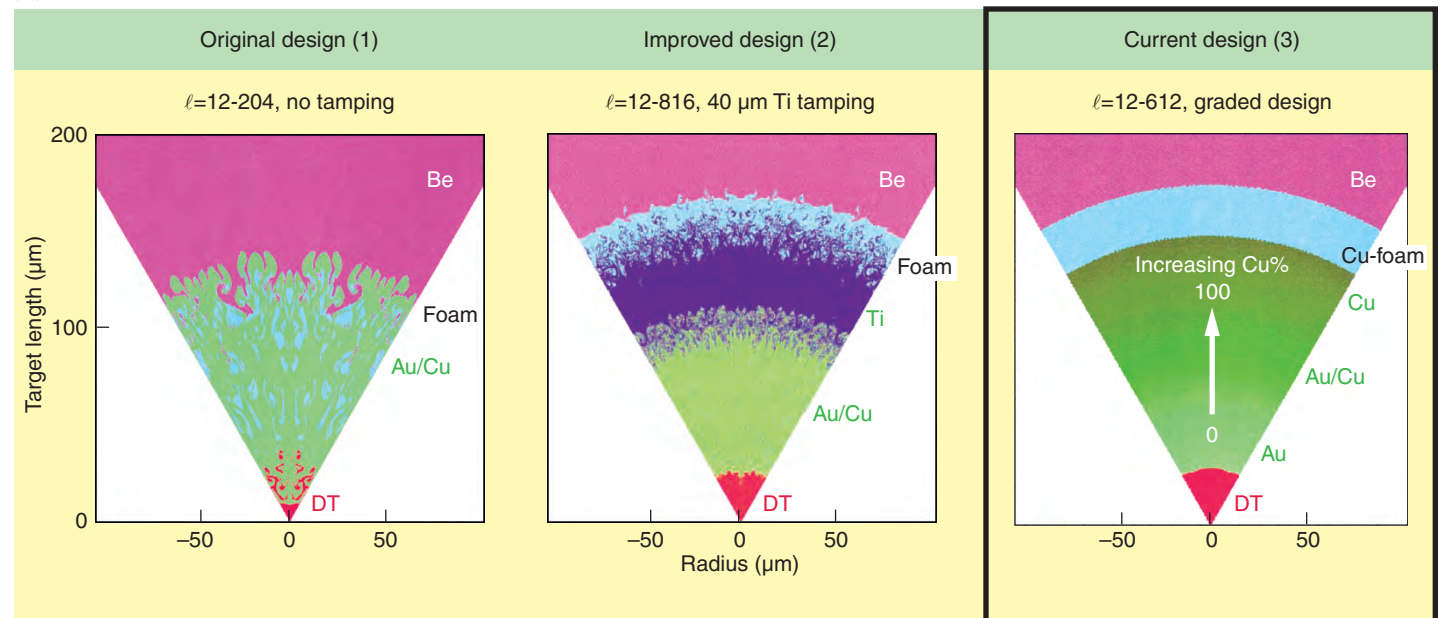
PROGRESS IN 2007

The current baseline ignition plan for NIF calls for the use of a cryogenic single-shell target capsule. Research has also been conducted on a complementary design, the double-shell (DS) target, which presents less demanding requirements for the laser power history and can be fielded at room temperature. Capsule surface roughness or asymmetries caused by the laser can

affect target performance. The surface roughness or asymmetries are particularly important for DS targets because they are highly sensitive to surface instabilities. Computation scientists used highly resolved HYDRA simulations to evaluate the viability and robustness of the DS and other novel target designs.


HYDRA simulations have been invaluable in guiding the design of DS targets. The original DS design proved to be highly susceptible to the growth of instabilities (figure a1). The current design, consisting of a graded bi-metallic shell and material-matching low-density metallic foam, overcomes the instability (figure a2)

(a)



RMS₁₂₋₆₀₀ = 3 nm

(a) Highly resolved multimode simulations at the time of maximum neutron production of (1) the originally proposed design, (2) a design that was modified by adding a 40-micrometer (μm) titanium (Ti) layer, and (3) an optimized design that uses a bi-metallic density-graded inner shell and low-density material-match metallic foam. Au = gold, Be = beryllium, Cu = copper, and DT = deuterium-tritium.



and a3). This work has generated significant interest at LLNL, culminating in a successful three-year Laboratory Directed Research and Development strategic initiative dedicated to developing these exotic nanomaterials.

HYDRA is massively parallel and scalable, and it runs on all Livermore Computing systems. The laser and radiation transport models use a novel parallelization technique optimized for load balance by multithreading. A genetic algorithm dynamically redistributes work for the most demanding phases of the computation.


The HYDRA development team continues to enhance the code's capabilities in response to user requests. A new physics module, for

example, will treat self-generated magnetic fields in 3D resistive magnetohydrodynamics (MHD), enabling unprecedented completeness in physics modeling and further improving scientists' understanding of target physics. A goal is for the new physics module to do the best possible job calculating the electron temperature field in the hohlraum, which allows a more precise evaluation of laser backscatter, radiation drive asymmetry, and laser-plasma interaction near the hohlraum entrance. Designers at Sandia and Lawrence Livermore national laboratories will make early use of this MHD module by simulating Z-pinch implosions on the Z-machine experiment.

SIGNIFICANCE

As LLNL scientists prepare for the NIF ignition campaign, HYDRA simulations are playing an important role in target design by simulating the laser-driven implosion characteristics of various designs. During the ignition campaign, it will be important to model ongoing experiments concurrently. Tuning experimental parameters to achieve optimal results requires an understanding of diagnostic results, aided by simulations run on a timely basis. The development of HYDRA continues to emphasize accuracy and speed to meet these requirements.

(b)



(b) A HYDRA simulation of a Z-pinch implosion in which a large azimuthal magnetic field compresses a plasma. The plasma is eventually compressed to a fraction of its original radius.

return
to TOC

3.03 — New Maxwell Solver Enables Scalable Electromagnetic Simulations

CONTACT INFORMATION


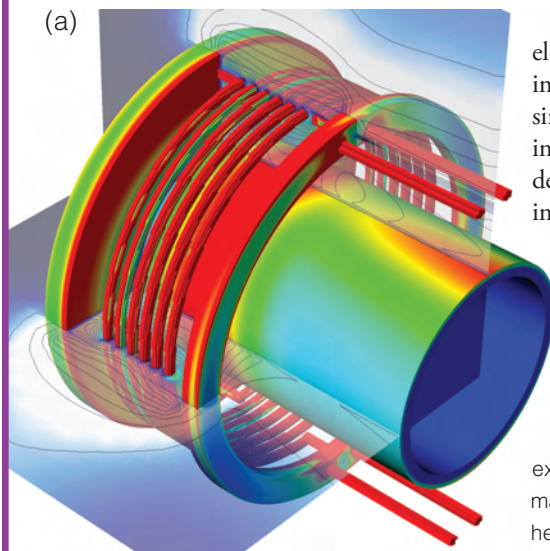
Tzanio Kolev, tzanio@llnl.gov
 Panayot Vassilevski, panayot@llnl.gov
 Web site: computation.llnl.gov/casc/hyprc



OVERVIEW

Numerical simulations of electromagnetic phenomena are critical in many practical applications. Such simulations often use Maxwell's equations in their second-order definite and semi-definite form. Their discretization results in large linear systems that are very challenging to solve. The Scalable Linear

(a) The new Maxwell solver speeds up a variety of challenging electromagnetic simulations. One example is the calculation of hydrodynamic stresses caused by large currents in pulsed-power experiments. The plot shows the transient magnetic field and eddy currents occurring in a helical coil with two side-by-side wires.



(b) Weak scaling results (73,000 unknowns per processor) for AMS, used as a preconditioner, and diagonally scaled conjugate gradient (DS-CG) applied to a simulation of electromagnetic diffusion in a coaxial cylinder. Using AMS is more than 25 times faster on the largest problem (73 million unknowns).

Solvers (SLS) project team developed a new auxiliary-space Maxwell solver (AMS) in the *hyprc* software library. AMS is the first provably scalable solver for unstructured definite and semi-definite Maxwell problems. Based on theoretical results by collaborators at Pennsylvania State University and the Swiss Federal Institute of Technology (ETH) Zurich, as well as research conducted by the SLS team, the algorithm works remarkably well on problems with large jumps in material coefficients and can handle regions with zero conductivity.

PROGRESS IN 2007

AMS was integrated and tested in two LLNL physics codes: ALE3D and HYDRA. The integration was done within new MHD packages for the two codes. Before AMS, the definite Maxwell linear solvers in these packages severely limited the size of the time step, the magnitude of material discontinuities, and ultimately, the resolution of the entire simulation.

The SLS team tested the new solver on several application problems while

completing an Advanced Simulation and Computing Level 2 milestone for *hyprc*. The results demonstrated that, when a large number of processors are used, AMS scales with the problem size and is dramatically faster than the previous MHD solvers. Refinements to the algorithm led to an additional 38% improvement in the time to solution on 1,024 processors.

AMS became the default solver for all electromagnetic calculations in ALE3D and is now available as an option in the HYDRA MHD package. In a recent large-scale ALE3D calculation, AMS required 47 iterations per time step to converge, while diagonally scaled conjugate gradient (the previously available solution method) did not converge within 10,000 iterations. Thus, AMS resulted in significant savings in total simulation time.

The scalability of the new Maxwell solver is exciting to the research community. AMS was highlighted at several major international conferences in 2007.

SIGNIFICANCE

The resolution of large-scale electromagnetic diffusion simulations based on Maxwell's equations is frequently limited by slow linear solvers. AMS is the first algorithm to deliver substantial and scalable speed improvements for such problems. AMS outperforms the previous solvers in MHD applications by as much as a factor of 25 and enables larger electromagnetic simulations than ever before possible.

return
to TOC

3.04 — Proteus Offers Timely Analysis in the Intelligence Arena

CONTACT INFORMATION

W. Marcus Miller
Kyle Halliday, halliday1@llnl.gov



OVERVIEW

In 2005, the National Intelligence Strategy for the United States established an objective to develop ways to analyze the most difficult intelligence targets. This goal can only be achieved if analysts exploit groundbreaking scientific advances in high-performance computing (HPC). Current technology makes it impossible for analysts to sort through the many petabytes of historical data and handle the daily deluge of new data. Traditional intelligence production systems cannot process data quickly or adaptively enough to keep up with subtle threat variations.

The Proteus Project (Phase I), a partnership between the Office of Naval Intelligence and LLNL, will deliver unprecedented computational power to analysts on the front lines, allowing them to rapidly explore spatial, temporal, and semantic relationships within a very large volume of data and to ask questions that would otherwise be computationally intractable.

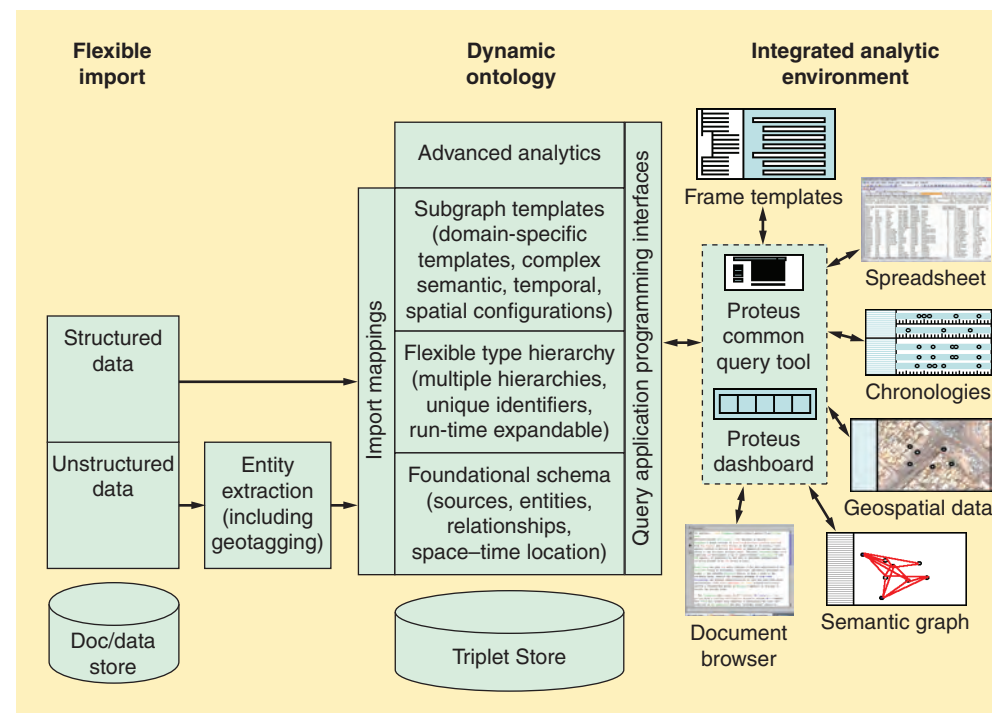
PROGRESS IN 2007

Proteus provides an analyst-centric approach that integrates HPC technology with multidimensional intelligence fusion software. Either structured (database) or unstructured (free-form message traffic and text) data can be ingested into the Proteus Triplet Store. In the Triplet Store, the information is organized in a generalized set of n -tuples that includes the triplet (entity, information source, geolocation) and other attributes as required. Once the

data are ingested, analysts can search for temporal, spatial, and semantic—relational connections. One requirement of the Proteus prototype design was that it use open-source codes, open standards, and published, service-oriented interfaces.

Proteus processing occurs in two major phases: Triplet-Store population and analysis against a dynamic ontology. The Triplet-Store population is run as a software pipeline comprising data ingestion, entity extraction and tagging, and Triplet-Store injection. The prototype system employs a multiple-terabyte storage capacity, shared-memory processor as a “landing zone” for incoming data. Data acquisition and entity extraction are performed using a 24-node IBM H20 blade center cluster. Results are then sent to a 32-node Netezza Data Analytics system for Triplet-Store injection and query processing.

The Netezza engine is a massively parallel Structured Query Language (SQL) system that processes terabytes of data in a fraction of the time required by a traditional database engine. Data are uniformly distributed on each of the 32 Netezza special processing units (SPUs), and SQL queries are processed in parallel by each SPU. Each SPU uses field-programmable gate array hardware to filter data to and from the disk units in real time. Queries are formulated in SPARQL, a run-time graph query language. Analysts interface with Proteus through the frame query tool, a client application that provides highly portable customization of the most common queries issued by an analyst.



The Proteus software architecture. Proteus extracts entity relationships from structured and unstructured data, allowing the analyst to discover correlations in seemingly unrelated events.

SIGNIFICANCE

Improving intelligence and information processing capabilities is essential for meeting 21st century security challenges. Using the Proteus system prototype, LLNL scientists applied state-of-the-art HPC technology to the extraction and analysis of huge volumes of intelligence data. This demonstration, along with Proteus's open-source implementation, will enable a broad range of groundbreaking results in intelligence analysis.



3.05 — Climate Modeling Research and Enabling Technology Have Worldwide Impact

OVERVIEW

LLNL is engaged in fulfilling the nation's need for climate modeling both by conducting its own research program and by providing enabling technology for researchers throughout the world. Research areas include high-resolution modeling, grand challenge scale computations on some of the most powerful computers in existence, regional climate modeling, coupled climate and carbon-cycle modeling, land-biosphere modeling, atmospheric chemistry modeling, detection and attribution of climate change, and model intercomparison projects. Scientists in Computation work closely with those in the Chemistry, Materials, Earth,

and Life Sciences (CMELS) Directorate to further these research efforts. Computation personnel and resources, under the aegis of the Program for Climate Model Diagnosis and Intercomparison (PCMDI), also play a crucial role in providing capabilities for scientists to manage, analyze, and visualize petabytes of data produced worldwide. Most notably, Computation scientists, in conjunction with CMELS scientists, managed and disseminated the Climate Model Intercomparison Project, Phase 3 (CMIP3) multimodel archive—generated by 23 distinct models—which was used for the Intergovernmental Panel on Climate Change (IPCC) Fourth Assessment Report.

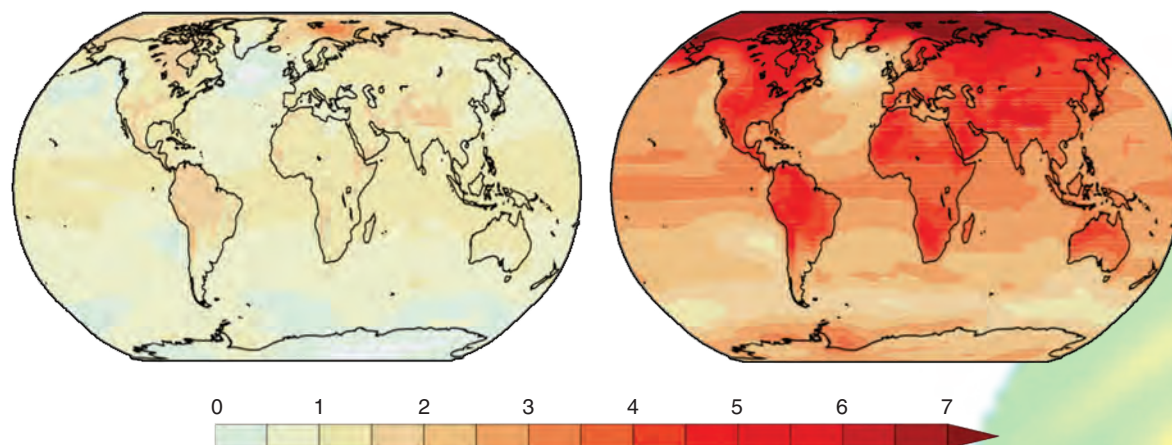
PROGRESS IN 2007

As part of the Scientific Discovery through Advanced Computing (SciDAC) Program, LLNL collaborated with other Department of Energy (DOE) labs and the National Center for Atmospheric Research (NCAR) to develop a comprehensive earth systems model. One of LLNL's primary roles in this collaboration was in the area of scalability. LLNL added a capability for various phases of the atmospheric calculation to proceed with different-sized domain decompositions, thereby enabling each piece of the computation to scale to its fullest potential. LLNL also provided

the capability for different phases of the time integration to proceed concurrently. These advances are particularly important for scenarios that contain computationally intensive model physics, such as interactive atmospheric chemistry, where the Laboratory has a prominent role in the collaboration.

Building on recent groundbreaking, multicentury, high-resolution coupled climate simulations on the Thunder computer (see Section 5.02), LLNL scientists extended the Community Climate System Model to enable studies on the Atlas computer. These studies will culminate in a multidecadal simulation at five times the

Average surface temperature change (°C) relative to preindustrial period, for early (left) and late (right) 21st century, as predicted by the 23-member CMIP3 archive.





CONTACT INFORMATION

Arthur Mirin, mirin@llnl.gov
 Dean N. Williams, williams13@llnl.gov


previous atmospheric resolution and 10 times the previous oceanic resolution. In a collaborative effort with Scripps Institution of Oceanography, Los Alamos National Laboratory, and NCAR, LLNL scientists demonstrated the capability for each component and various subsets to run properly and effectively at this unprecedented resolution and commenced fully coupled calculations.

In a study to detect and attribute causes of regional climate change, results of LLNL's global simulations were used as boundary conditions for an even higher resolution regional simulation focused only on the western United States. Combining regional calculations with statistical analyses of output datasets, LLNL scientists worked in collaboration with Scripps to analyze changes in

temperature, hydrology, snowpack, and stream-flow timing, and the extent to which these changes are human-induced. Articles in these areas were submitted to leading journals.

Enabling climate scientists to efficiently access, manipulate, and analyze data is an overarching problem. Researchers are increasingly required to access large, complex datasets that are archived in various formats on disparate platforms scattered throughout the world. SciDAC funds a multilaboratory consortium known as the Earth System Grid to

advance climate science by developing computational resources and technologies that researchers can use to access and manage model data that are physically located in distributed multiplatform archives—worldwide. With LLNL as the lead institution, the Earth System Grid currently manages the CMIP3 multimodel (35-terabyte) archive, which has distributed more than 350 terabytes of data to 1,650 registered analysis projects. More than 400 peer-reviewed publications have been authored on analyses of the CMIP3 data.



A winter storm moving eastward across California produced from a 6-kilometer-resolution simulation using an established regional climate model. Lower panel shows cloud concentration; upper panel shows surface precipitation rate, ranging from 2 millimeters per hour (green) to 20 millimeters per hour (red).

SIGNIFICANCE

Computation, through its partnership with CMELS, plays a vital role in maintaining the Laboratory's position at the forefront of climate simulation. Such efforts were recognized in the collective award of the 2007 Nobel Peace Prize to the IPCC scientific community (shared with former U.S. Vice President Al Gore). Areas of high impact include advanced model development, high-resolution modeling, and scaling to leadership computing systems as well as data management and analysis capability as provided through PCMDI. The ultrahigh-resolution coupled simulations planned for Atlas will be unprecedented in terms of resolution and simulation time. Community efforts have already begun to extend this capability to BlueGene-class architectures. Key to the success of these efforts is collaboration with scientists throughout the climate-change community. LLNL and Computation are positioned to play a leading role in the next phase of climate studies that will provide guidance in addressing the effects of and mitigation strategies for global warming.



3.06 — Cyber Security Research Creates Innovative Tools to Protect Computer Networks

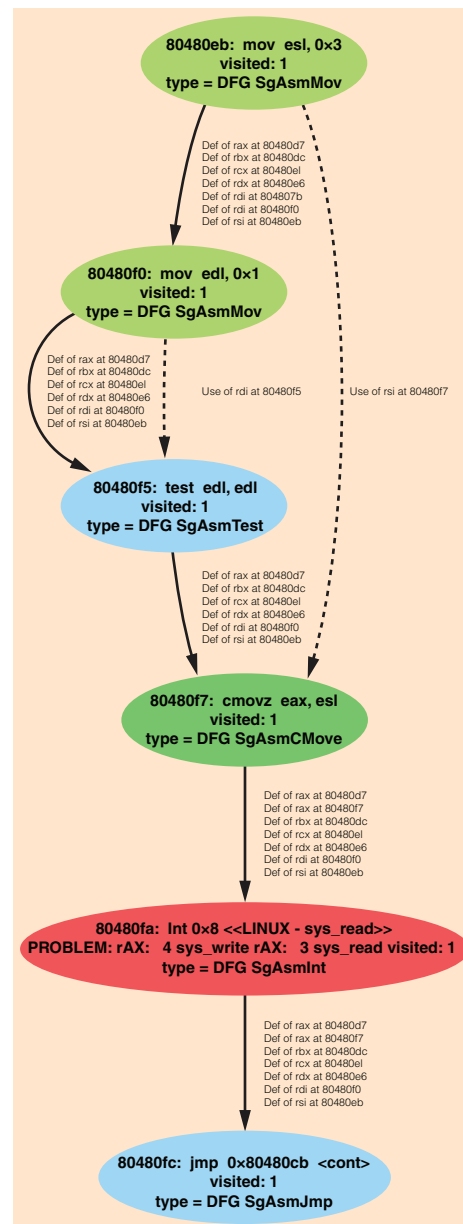
CONTACT INFORMATION Dan Quinlan, dquinlan@llnl.gov



OVERVIEW

Cyber security is a broad and increasingly important research topic with far-reaching implications to LLNL and its customers. One topic Laboratory researchers are investigating is the automated detection of security flaws and malicious code in software. DOE obtains a significant amount of software from external vendors, either as purchased binaries or from open-source software suppliers, and LLNL produces large quantities of software for both internal use and for external customers, such as DoD and DHS. Security analysis of software is an important step in strengthening DOE and the Laboratory's cyber security defenses and in providing tools to support secure software development. Building on computer program analysis and optimization techniques developed in the ROSE project, which is funded by the DOE Office of Science, LLNL scientists are focusing their efforts on security analysis and the handling of binaries.

The control flow and data flow over six instructions from a large binary executable is shown graphically. Visualization techniques classify the instructions and assign colors. The red instructions in the graph are identified as a write system call.



PROGRESS IN 2007

LLNL research is being compiled into a tool called Compass, which is built using the ROSE software. Compass automates security analysis of source code in C, C++, and Fortran 2003. The Compass package is being developed in collaboration with Carnegie Mellon University's Computer Emergency Response Team, the National Institute of Standards and Technology, and other external organizations. Compass is uniquely extensible and permits arbitrary security rules to be defined while automating the detection of rule violations in large-scale source-code applications.

LLNL researchers have anticipated the aggressive security analysis technology that is in development and have pioneered novel

parallel attribute evaluation techniques for shared- and distributed-memory parallel computers. Parallel attribute evaluation is at the center of LLNL's security research for large-scale parallel supercomputers. One goal is to make expensive forms of program analysis, not tractable on serial computers, run in parallel on large-scale computers. Binary analysis permits disassembly and examination of large-scale applications where no source code is available. This work also benefits from the parallel program analysis infrastructure. The identification of system calls in the binary is an important step in deducing the behavior of an arbitrary binary executable. In 2007, LLNL cyber security researchers developed significant internal and external collaborations to pursue these lines of inquiry.

SIGNIFICANCE

Fundamentally, a computer network attached to the Internet is only as secure as the software that runs on its computers. In today's world, where adversaries regularly exploit the vulnerabilities of computer networks by attacking the software programs that run on them, cyber security research takes on increasing importance. Open-source and third-party binary software is routinely used both inside and outside the government. But as software cannot generally be trusted, the burden of verifying its security falls to the users. LLNL's research will provide users with tools for analysis and security verification of software.



4

Computational Scientists Facilitate Programmatic Success

With approximately 500 computational scientists, including researchers, applications programmers, and discipline experts, the Computing Applications and Research Department in the Computation Directorate is one of the largest applied computer science organizations in the world. These scientists partner with every directorate, and nearly every program, at Lawrence Livermore National Laboratory (LLNL). One of their most important tasks is to develop software that enables programmatic applications scientists to fulfill the mission requirements of LLNL's sponsors in the National Nuclear Security Administration (NNSA), the Department of Energy (DOE), the Department of Homeland Security (DHS), the Department of Defense (DoD), and other federal and state agencies.

Computation personnel are “matrixed” to the programs—they work daily with the programmatic project teams and are often in the assignment for many years. The projects range from small two- or three-person teams to very large teams of 20 or more scientists. As a result, Computation personnel develop deep domain expertise in a far-reaching set of fields, including materials modeling, fluid dynamics, genomic sequencing, climatology, geophysics, astrophysics, scalable algorithms, parallel processing, and real-time control systems. This

4.00 — Computational Scientists Facilitate Programmatic Success

CONTACT INFORMATION Patricia Damkroger, damkroger1@llnl.gov



extraordinarily broad “corporate” expertise enables Computation to respond readily to sudden or dramatic shifts in the needs of the Laboratory. The stories in this section highlight some of the critical work Computation does for its programmatic sponsors and focus on how the successes of 2007 pave the way for the future.

This section begins with a set of articles describing the role that Computation has played in creating software frameworks used to develop the large-scale Integrated Computer Control System (ICCS) for the National Ignition Facility (NIF). NIF’s 192 beams are partitioned into 24 independent eight-beam bundles, each containing laser hardware that is modularized into line-replaceable units such as optical assemblies, laser amplifiers, and multifunction sensor packages. The whole system is operated by ICCS in a distributed architecture using more than 800 front-end processors and 50 supervisory servers. ICCS controls the beams using 60,000 control and diagnostic points. NIF’s automated control subsystems are built from a common object-oriented software framework that deploys the software across the computer network and achieves interoperation between different languages and target architectures. In 2006, Computation personnel deployed a shot automation framework to orchestrate and automate shots. Only one year later, in October 2007, under the control of ICCS, 96 beams were fired simultaneously, demonstrating that the independent bundle control system will scale to the full 192 beams. By year’s end, 120 beams (15 bundles) had been commissioned and had demonstrated 1.4-megajoule capability of infrared light.

Computation employees are also helping prepare NIF for completion in 2009. They are heavily involved in the development of automatic alignment and target positioner systems, laser performance models, and the data repository systems that will support NIF experiments. Additional capabilities to support fusion ignition shots in a National Ignition Campaign beginning in 2010 will include a cryogenic target system, target diagnostics, and integrated experimental shot data analysis with tools for data visualization and archiving.

Computation’s applications programmers support many other programs as well, and several of the stories that follow highlight important achievements elsewhere at the Laboratory. For example, Computation programmers and scientists hold critical positions in the Global Security organization. Computation’s role in pathogen bioinformatics far exceeds a mere supporting role, and Computation employees drive results in biodefense. The directorate is also a key player in the successful Context-Aware Nuclear Evaluation System project, which is devoted to improving radiation detector performance by incorporating contextual information into the analysis of detection data.

Traditionally, Computation employees have been closely linked with the mission of the Weapons and Complex Integration Principal Directorate. More than 100 applications developers currently support LLNL’s work in this area. Computation has a long-term involvement in numerous multiphysics, multidimensional codes, such as ALE3D, ARIES, KULL, and others. ALE3D is an extremely large code supporting many users throughout the Laboratory, the NNSA complex, and the general computational community. ALE3D, like all large multiphysics packages, requires talented computer scientists to maintain an agile yet structured code base. A dedicated team from Computation are key to the success of ALE3D.

In addition to supporting the main programs at the Laboratory, Computation employees serve the broader weapons complex. The Argus system, for example, protects high-security facilities, special nuclear materials, and Top Secret and Secret Restricted Data. This system comprises a computing kernel that controls access points through portals. Argus has been evolving for the past 20 years and continues to upgrade with new technologies and security requirements. Computation programmers guide and facilitate this evolution.

Considering the breadth of Computation’s involvement with the Laboratory programs, no small set of highlights can possibly capture the overall impact of the computer scientists, applications programmers, and discipline scientists that the directorate deploys. The stories in this section serve merely as a sample of how the Laboratory is served by these talented and dedicated people.



4.01 — NIF Integrated Computer Control System Software Demonstrates Full-Scale Operation

CONTACT INFORMATION Robert Carey, carey3@llnl.gov



OVERVIEW

NIF is operated by ICCS, which is a scalable, framework-based control system distributed over hundreds of computers throughout NIF's 192 laser beam experimental-physics machine. The framework provides templates and services at multiple levels of abstraction for the construction of software applications that communicate via the Common Object Request Broker Architecture (CORBA). Object-oriented software design patterns are implemented as templates and extended by application software. Developers extend the framework base classes to model the numerous physical control points and implement specializations of common application behaviors. At full scale, approximately 140,000 software objects are active, each individually addressable through CORBA and distributed throughout 1,500 processes on more than 800 computers.

PROGRESS IN 2007

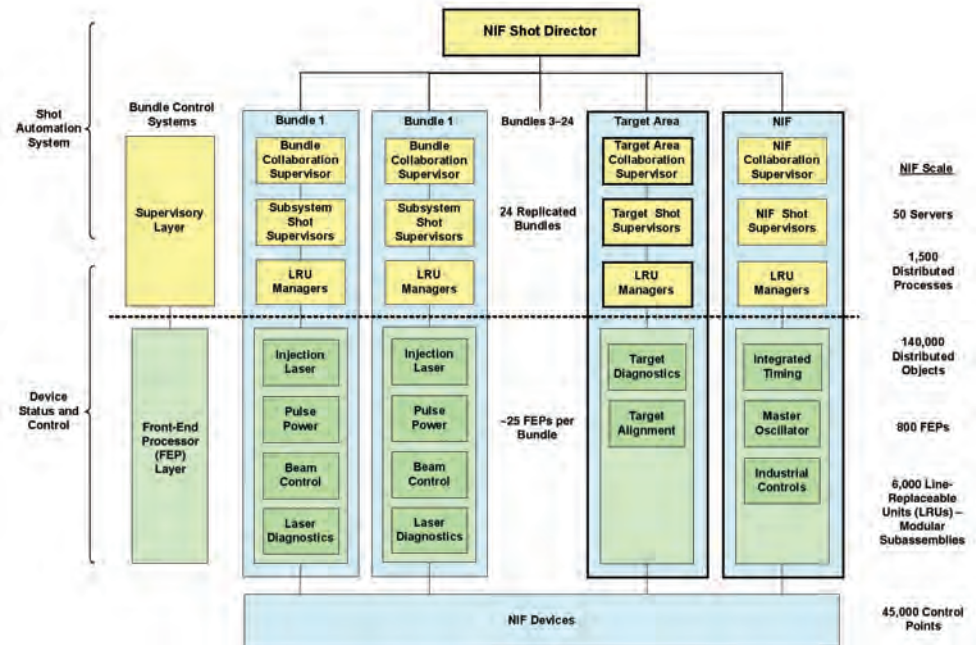
ICCS applications are developed using Ada95 or Java, CORBA, and object-oriented programming. Ada95 is used to implement most of the control system semantics. Java is used to produce graphical user interfaces and some of the central services. CORBA provides location- and language-transparent distributed communication using Transmission Control Protocol/Internet Protocol transport.

The control system computers and processes are largely partitioned by the hardware and software objects required

to operate one NIF bundle (eight laser beams), with the exception of the central services and applications controlling front-end, target area, and industrial controls systems, which are not replicated. This partitioning is a key element in ensuring the scalability of the control system.

The top layer of the ICCS software architecture is the shot automation system. This software provides a highly flexible framework for performing NIF shots and manages the progress of a NIF experiment through a sequence of states. In each state, partitioned collections of processes manage the workflow in a sequence of steps defined by a shot model, which is represented by a large directed graph. Many thousands of actuators, instruments, and sensors are automatically configured, verified, and operated reliably in parallel by the data-driven shot automation system. Laser experiments using all 192 beams were conducted in 2007 to demonstrate control system operation at full scale. Target area subsystems are being activated in preparation for target experiments in 2009 and for fusion ignition experiments in 2010.

In 2007, a team of LLNL computer scientists began migrating and refactoring the Ada95 codes to Java. This improvement positions the controls architecture for better long-term survivability and leverages substantial industry support and expertise in Java applications. Using CORBA makes this migration seamless in terms of platform interoperability and language transparency.



Hardware and software partitioning by bundle is key to ensuring the scalability of the control system.

SIGNIFICANCE

Full-scale operation and automation of NIF is essential to meet NIF Program missions. In 2007, the ICCS demonstrated this significant capability. Technological innovation and evolution in the computer industry are influencing the migration of ICCS languages and software development tools toward the mainstream. The ICCS software architecture and framework have shown the flexibility to support this technological evolution without significant re-engineering. Maintaining pace with the computer industry is crucial for the long-term operation and support of NIF.



4.02 — NIF Automatic Alignment Shows Accuracy, Reliability, and Speed

CONTACT INFORMATION Ben Horowitz, horowitz2@llnl.gov




OVERVIEW

The ICCS automatic alignment system for NIF is a software control system that concurrently directs NIF's 192 laser beams along a 300-meter beam path to a 50-micron focus at the target chamber center. This feat is accomplished in less than 30 minutes. In 2007, automatic alignment was used successfully in 500 high-energy system and 1,500 laser qualification shots.

The automatic alignment system contains control systems for each of NIF's 24 eight-beam bundles and the target area.

Each beamline contains three segments (preamplifier, main laser, and target area) and is controlled in largely independent fashion by two operators. Each segment contains control loops that use image processing (to determine the beam location) and adjust motors to move the beam to its desired position. Approximately 9,000 motors, 3,000 actuators, and 150 digital cameras are used by the system. The scale of the system requires reliable and efficient coordination of shared devices, such as sensors that monitor multiple beams.

(a) A beam control operator supervises the automatic alignment system. A single operator supervises the alignment of the main laser and preamplifier module.



(b) The target positioner and alignment system precisely locate a target in the target chamber. The target is positioned with an accuracy of less than the thickness of a human hair.

PROGRESS IN 2007

A major focus of effort in 2007 was development of the component mediation system (CMS). CMS coordinates device movements and reserves devices for exclusive or shared use. CMS allows control loops to perform in parallel whenever possible and to perform serially when necessary. It manages shared devices using an object called a mediated component (MC) that provides concurrency-controlled commands. MCs coordinate client requests to shared devices by managing a queue of active requests. Access to an MC object that is already reserved is blocked until the MC is available. Deadlocks are eliminated by enforcing a priority order for using MC

objects. Waiting for device movements dominates overall alignment time. To optimize performance, clients waiting for the same resource move to the front of an MC object's queue, yielding fewer device movements and increasing performance.

A second focus of effort in 2007 was improved detection of off-normal beam-location images. A unique classification algorithm for NIF images was implemented, where images are sub-sampled with spiral search patterns, clearly identifying edge reflections on corner cubes (three mirrors mounted at right angles) in the NIF beamlines. This process continues to provide a high level of confidence for normal/off-normal classification.

SIGNIFICANCE

The automatic alignment system meets requirements for accuracy, reliability, and speed. The novel signal and image processing techniques developed at LLNL enable NIF's image processing to accurately and reliably estimate image centers. CMS efficiently coordinates shared devices. The automatic alignment system is exercised an average of four times per day on the 15 currently operational bundles. The ICCS bundle-based architecture allows the automatic alignment system to scale to the full NIF configuration of 24 bundles and to support alignment to the target chamber for ignition experiments beginning in 2010.



4.03 — Automated Target Positioners Reduce the Setup Time for NIF Shots

CONTACT INFORMATION Ephraim Tekle, tekle1@llnl.gov



OVERVIEW


The NIF target chamber is a 10-meter-diameter sphere under vacuum environment with a thick protective wall that is penetrated by the arriving laser beams and 10 large target and diagnostic instrument positioners. Target chamber shot experiments require a variety of positioner arrangements near the chamber center to be aligned to an accuracy of 10 micrometers. Positioners are some of the largest devices in NIF, and they require careful monitoring and control in three dimensions to prevent interference while being maneuvered within a few millimeters of each other. The NIF ICCS provides efficient, reliable, and flexible multipositioner controls, which are implemented using advanced video-control integration that incorporates remote position sensing, real-time visualization and analysis of a computer-aided design model of target chamber devices, and automated alignment capabilities.

PROGRESS IN 2007

For a NIF shot to be fired, all positioners must be aligned, and the 192 laser beams must be precisely aimed at the 10-micrometer target. Positioner alignment is a time-consuming procedure if done manually. Automation of the chamber alignment process helps shorten the shot cycle. This automation is supported by multipositioner move sequencing, automated path finding, and model-based real-time analysis and visualization tools.

Operators monitor positioner movements using the Chamber Interior Viewing System (CIVS) video cameras. CIVS magnifies the central region of the chamber to permit precise adjustments and, thus, limits the camera field of view enough to reduce situational awareness for the operators. To mitigate this problem, operators can move positioners along virtual axes that are defined in the field of view of each CIVS camera. Cross-coupling matrices ensure the positioner drives follow these virtual axes while providing depth boundaries, speed limits, and motion increment restrictions. Calibrated video views allow operators to “drag” a movement command across live video of the positioner. The control system translates this command into a complex set of motor moves and describes the proposed move to the operator in human terms, indicating which motors will be used for the move. Once the move is approved, the positioner is placed in motion, and the operator observes the motion, tracking the initial “dragged” command in real time.

A dedicated test laboratory integrates positioner simulators, digital cameras, programmable logic controllers, and other equipment. Hardware-based testing provides a high-fidelity offline test environment. Software-based testing uses emulated virtual positioners to allow multiple tests to be conducted concurrently.



Screen shot from a driving video of a target alignment positioner control panel. The alignment sensor positioner is being moved from its current position to its final alignment position. The overlay (green) aids operators in dragging (yellow arrow) the target alignment sensor positioner to its final position. The live video feed is calibrated (yellow coordinate lines) and move boundaries are shown (red line).

SIGNIFICANCE

The model-based visual control system is an advanced integrated operations tool for coordinating the complex alignment of NIF positioners where a large fraction of the chamber interior is inaccessible for viewing. Model validation ensures the virtual system accurately represents the real system. Model-augmented live videos provide feedback to the operator during manual control operations. The automated controls keep the positioners interference-free and provide operators with a view of the chamber interior from any vantage point at all times.

4.04 — NIF's Laser Performance Operation Model Facilitates End-to-End Shot Planning and Analysis

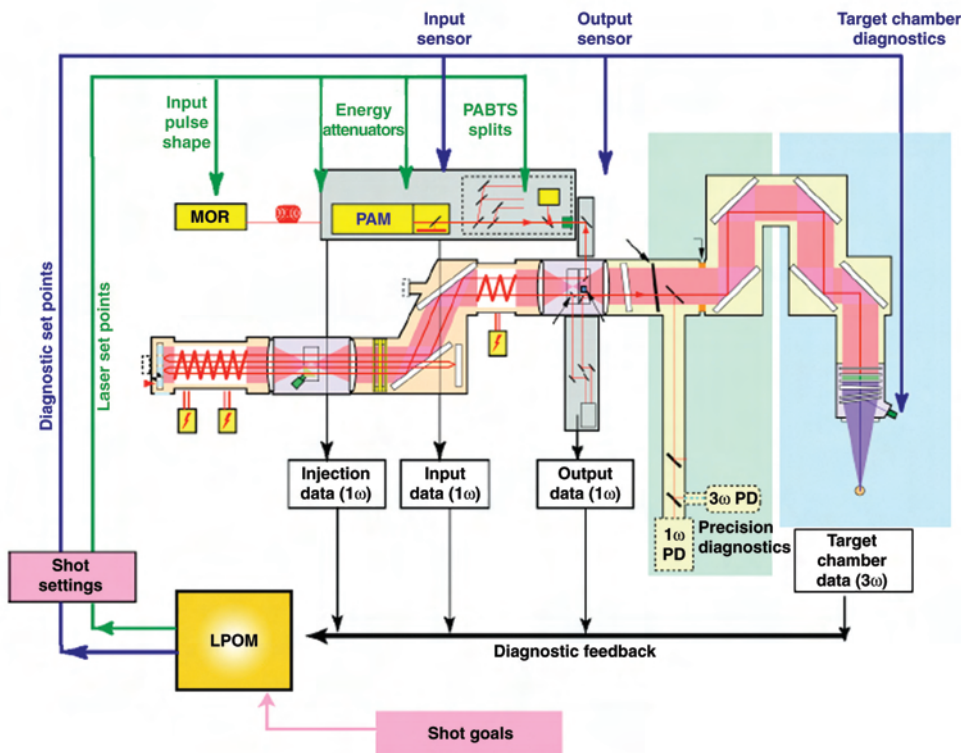
CONTACT INFORMATION Ronald House, house6@llnl.gov



OVERVIEW

The Laser Performance Operations Model (LPOM) is a complex combination of computing hardware and software developed to model and predict laser energetics of NIF and to automate the determination of setup parameters of the laser to meet the goals of NIF experiments. Prior to each NIF shot, LPOM determines the characteristics of the system required to achieve the desired

main laser output, validates the settings for equipment protection, and defines the required diagnostic setup. After each NIF shot, LPOM works within the ICCS to supply NIF operators with post-shot data analysis and recommend adjustments to settings where appropriate. LPOM was deployed prior to NIF's first main laser shots and has since been used to set up and analyze every shot.



LPOM automates shot setup and uses post-shot analysis to maintain an accurate model of NIF.

PROGRESS IN 2007

Ongoing commissioning activities, along with several intensive precision diagnostic system campaigns, were used in 2007 to demonstrate the scalability of LPOM and its ability to support the increasing demands of NIF operations. This process tested both the LPOM software and the hardware of the LPOM cluster, currently comprising four interactive nodes and 96 compute nodes. Early in the year, the LPOM cluster was connected to the ICCS Test Bed, allowing ICCS to test multiple-bundle experiments and shots.

During the ICCS 5.4 shot regression testing, the first end-to-end 192-beamline setup calculation was successfully demonstrated, and before the end of the year, the first full NIF system shots (192 beams) were verified by LPOM. For these shots, LPOM successfully read all diagnostics data and analyzed laser performance for the complete system in approximately four minutes.

LPOM developers have been leaders in creating Web-based visualization tools for NIF. During 2007, they extended the LPOM visualization functionality

by improving its Web response time and enhancing its interactive trending tools. Developers also added many features to LPOM's Web services, including advanced data integrity alerting and reporting, full cluster overviews of energetics, power traces and images, experimental campaign summaries, LPOM model status, improved shot setup reporting, LPOM and database performance statistics, and LPOM storage utilization.


Beyond providing support for daily shot operations, the LPOM team continues to develop additional functionality required to support future NIF activities. This work includes implementing full NIF final optics assembly geometry to support calculations of the 192-beam target chamber center energetics, automation of a 1-watt pulse correction algorithm, implementation of a new power balance model interface, support for faster multiple cluster shot verifications, improved LPOM server performance for multiple concurrent connections, automated test suite for post-shot analysis regression testing, and reduced storage requirements.

SIGNIFICANCE

Success on many of NIF's missions depends on obtaining precisely specified high-power waveforms from each of the 192 laser beams over a variety of pulse lengths and temporal shapes with the minimum possible optical damage. LPOM is designed to help achieve these goals while simultaneously supporting NIF shot planning, operations, and reviews.



4.05 — NIF Data Repository Manages Long-Term Storage and Provides Analysis and Visualization Tools



CONTACT INFORMATION

John Rhodes, rhodes1@llnl.gov

Timothy Frazier (NIF), frazier4@llnl.gov

OVERVIEW

The NIF Data Repository (NDR) is a petabyte-scale database that provides long-term storage of datasets resulting from experimental campaigns run on NIF. At its heart, the NDR is a highly parallelized database that captures several types of data for each NIF experiment: identities of the team leaders and members, goals of the experiment and technical setup parameters, configurations and calibrations of the thousands of components that make up NIF and its diagnostic systems, raw experimental results for target and laser system experiments, and processed results that are computed by the thousands of analysis modules that automatically run following each shot.

PROGRESS IN 2007

NIF's 30-year lifetime requires a comprehensive approach for capturing, managing, analyzing, and visualizing the many petabytes of data that NIF will generate. In addition to ensuring that large volumes of data remain accessible well beyond NIF's operational lifetime, the NDR must provide an easy-to-follow audit trail for all calculated results, a concept known as data pedigree management (DPM). DPM capabilities were added in 2007 to provide experimentalists with explicit knowledge of the data used in each calculation, from setup values (such as camera sweep speed settings) to the versions of the analysis programs that computed each result.

Furthermore, should there be a change in any data item or analysis code used for a calculation, DPM enables the automatic rerunning of the programs needed to compute the result.

Three analysis suites for target diagnostic cameras were installed in 2007. Another 10 to 15 diagnostics, with analysis suites of their own, will follow in 2008. Also in 2008, NDR analysis programs will be used to calculate results for virtual experimental campaigns run on NNSA's Advanced Simulation and Computing supercomputers.

The NDR also supports the visualization of datasets. With NIF producing prodigious quantities of information, the ability to quickly find and review both raw datasets and calculated results is a critical capability in a number of contexts: immediately after an experiment to assess results (and possibly to alter settings and goals for the next experiment), shortly after a series of experiments to make decisions about the scientific merit of upcoming campaigns, and indefinitely after an experiment to validate and publish results and conclusions for the scientific community.



SIGNIFICANCE

The NDR will provide experimentalists and collaborators from the worldwide fusion research community with a single source for the datasets generated by NIF experiments. The design and architecture of the NDR, including pedigree management and data visualization capabilities, will provide for the efficient management of datasets whose structure is currently well understood and will provide the platform for those whose structure and significance come to light in the future.

Data processing analysis and visualization tools for on-line optics inspection determine optics replacement needs.



4.06 — Computation Provides Leadership to Handle Next-Generation Biothreats

OVERVIEW

In 2000, LLNL computer scientists were the first to suggest using whole-genome comparative analyses techniques to automate the design of DNA-based signatures that detect pathogens. This idea became a central process in the Laboratory’s KPATH pipeline, a fully automated system to design and maintain DNA-based signatures. Developed in collaboration with the Centers for Disease Control and Prevention, KPATH downloads new and updated genomic sequences from major public databases, compares these sequences with those already in use, and determines whether any new sequences have invalidated the existing ones. The KPATH pipeline has significantly increased the efficiency of pathogen detection. As a result, LLNL has become a national leader in pathogen detection technology. Many of the DNA signatures used in the Biological Aerosol Sentry and Information System and BioWatch System, which provide dozens of U.S. cities with 24/7 protection, were developed using this automated process.

However, identifying a pathogen’s species does not provide sufficient information for public health authorities to take action. When emergency response personnel determine that a detected organism is a near neighbor to a high threat, they must quickly determine

KEY

BA discriminating genes
BC discriminating genes
BT discriminating genes
BA virulence genes
BA strain discrimination
BT strain discrimination

whether this organism poses a threat to human health. The specter of genetic engineering, whether an act of nature or terrorism, brings another concern—that harmful mechanisms might suddenly appear in previously benign organisms or that bioterror threat organisms might begin to show resistance to the stockpiled antibiotics.

Virulence.group	Name	BA 0442	BA Sterne	BA 0193	BT 97_0027	BT AI Hakem	BC 7
toxin	L1 component of the B. cereus HBL tripartite haemolytic toxin	0	0	0	1	1	1
toxin-anthrax	lef, zinc metalloprotease that cleaves MAPKK	1	1	1	0	0	0
toxin-anthrax	cya, calcium/calmodulin-dependent adenylylcyclase	1	1	1	0	0	0
toxin-anthrax	CIII/pagA provides entry of B.anthraxis lef and cya into host cell	1	1	1	0	0	0
multi-drug efflux pump	putative major facilitator superfamily export protein	1	1	1	0	0	0
capsular polysaccharide synthesis	wzt2, bexA, putative ABC transporter for capsular polysaccharide	1	1	1	0	0	0
iron acquisition-siderophore synthesis	pchE, irp2, HMWP2 pyochelin/yersiniabactin synthetase	1	1	1	0	0	0
LPS biosynthesis	wzt, ABC transporter, ATP binding component	1	1	1	0	0	0
iron acquisition-siderophore synthesis	pchC, lipase similar to pyochelin biosynthetic protein from P.aeruginosa	0	0	0	0	0	1
antibiotic resistance	cpmF, carbapenem resistance	0	0	0	0	0	1
antibiotic resistance	BPSL2708, putative exported metallo-beta-lactamase	0	0	0	0	0	1
antibiotic resistance	borB bacitracin transport permease protein	1	1	1	1	1	0
antibiotic resistance	mphA like, macrolide 2-phosphotransferase I	1	1	1	1	1	0
nitrogen metabolism	aliphatic amidase (EC 3.5.1.4) (H.pylori amiE)	1	1	1	1	1	0
nitrogen metabolism	formamidase (EC 3.5.1.49) (H.pylori amiF)	1	1	1	1	1	0
antibiotic resistance	bsr blastidicin-S deaminase (EC 3.5.4.23)	1	1	1	0	0	1
capsule biosynthesis	capD, required for capsule biosynthesis in Bacillus anthracis	1	0	1	0	0	0
capsule biosynthesis	capB, capsule synthesis B. anthracis	1	0	1	0	0	0
capsule biosynthesis	capC, capsule synthesis B. anthracis	1	0	1	0	0	0
antibiotic resistance	borC bacitracin transport permease protein	1	0	1	1	1	1
antibiotic resistance	ksgA, erm(1,A,B,C) Dimethyladenosine transferase (EC 2.1.1.1)	1	1	0	1	1	1
antibiotic resistance	bla(2,B) Beta-lactamase IMP-1 precursor (EC 3.5.2.6)	1	1	0	1	1	1
TAT secretion	tatD, TAT secretion, sec independent protein translocase, de	1	1	0	1	1	1
transcriptional control	agrC, quorum-sensing sensor kinase	1	1	0	1	1	1
antibiotic resistance	mbcG resistance to the peptide antibiotic microcin B17	1	1	1	1	0	1
antibiotic resistance	bis blastidicin S-acetyltransferase (EC 2.3.1.-)	0	0	0	1	0	1
surface antigen synthesis	glycosyltransferase 2(F.tularensis FTT0797, FTT0798)	0	0	0	1	0	0

A Livermore-developed microarray probe identifies the presence (red) or absence (green) of virulent or antibiotic-resistant genes to differentiate the species and strain of six organisms examined in virulence-array experiments. BA stands for *Bacillus anthracis*, BT stands for *Bacillus thuringiensis*, and BC stands for *Bacillus cereus*.

To address these challenges, LLNL is creating a microarray to detect all known virulence and antibiotic-resistance mechanisms from relevant high-threat organisms. This project, which is funded by DHS, involves a multidisciplinary team of researchers from LLNL’s Microarray Center; the Chemistry, Materials, Earth, and Life Sciences (CMELS) Directorate; and Computation.

PROGRESS IN 2007

Microarray probes for this project comprise specific strings of 50 to 65 DNA nucleotides A/T/G/C selected to accurately discriminate the targeted mechanisms. Using LLNL’s Thunder and Yana machines, computer scientists extensively analyzed designs to ensure that the microarray probe detects and clusters all




the relevant mechanisms for the targeted genes. Scientists in CMELS manufactured the microarray and performed laboratory experiments to test its accuracy.

Experiments on the causative agents of anthrax, plague, and tularemia demonstrated that the microarray could provide a functional genetic “parts list” for an unknown sample to enable quick and accurate determination of the species and strain of an organism. Test results showed higher degrees of resolution for anthrax and plague than for tularemia because the precise mechanisms involved in tularemia are not completely understood.

When briefed on these results, DHS senior management expressed interest in using the microarray to help BioWatch characterize future low-level detections of uncharacterized organisms that are near neighbors to known threat pathogens. Rapid determination of the virulence and antibiotic-resistance “parts list” would provide public health officials with information for determining appropriate responses. Another application for the microarray is to delineate the presence-and-absence pattern of functionally important genes, providing a “functional fingerprint.” This new fingerprint will produce valuable information to complement data from the existing typing schemes used in microbial forensics.

In 2007, the Livermore team also designed and tested a microarray for the



Viral Detection Project (VDP), which is funded by the Laboratory Directed Research and Development Program. The VDP microarray chip can place an unknown virus into one of the 72 vertebrate-infecting viral families. This microarray is an updated version of the one used in 2003 by UC San Francisco to determine that SARS (severe acute respiratory syndrome) is a new type of coronavirus. Initial tests on the VDP chip indicate a high degree of success in the microarray probe design.

The viral discovery chip was tested on a set of blinded clinical samples provided by a UC San Francisco collaborator. Analysis of the data showed that the Livermore team correctly detected the presence of an adenovirus in that sample.

SIGNIFICANCE

Livermore is a national leader in bioinformatics, designing tools to deal with next-generation biothreats. In 2008, an LLNL research team will develop an extensive virulence array for the DHS microbial forensics center. Another team will supply DHS, the Defense Threat Reduction Agency, and the intelligence community with planning information on countering advanced biothreats. The VDP team, in collaboration with UC San Francisco and the University of Texas Medical Branch at Galveston, will test the viral discovery array to determine its ability to classify unknown viruses. As part of the VDP microarray effort, Computation scientists have, for the first time, acquired a wet lab and hired biologists to verify the computational predictions.

These projects highlight the critical role bioinformatics can play in applied research for biodefense. By leveraging LLNL supercomputers, this multidisciplinary team is applying the Laboratory's expertise to develop technologies that can detect known mechanisms and classify viruses.

4.07 — Argus Upgrades Safeguard National Security Facilities

CONTACT INFORMATION Kimberly Ferrari, ferrari2@llnl.gov



An Argus central alarm station.

OVERVIEW

Argus was originally developed in the late 1980s as part of an institutional effort to upgrade the Laboratory's physical security system. The sophisticated, computerized system was designed to monitor and control entry into the Laboratory's high-security buildings, particularly those that contain Category I/II Special Nuclear Material and material classified as Top Secret or Secret Restricted Data. Argus was designed to meet the stringent security requirements contained in DOE's security orders.

Since the early success of the Argus system at LLNL, the system has been deployed to the Joint National Integration

Center at Schriever Air Force Base in Colorado, Idaho National Engineering Laboratory, DOE Headquarters, Pantex, and, most recently, Los Alamos National Laboratory. In 2004, DOE cited Argus as the standard for physical security systems protecting facilities where the consequences of intrusion are significant. Plans are under way to implement Argus at the Savannah River Site and Y-12 National Security Complex.

The Argus team comprises personnel from the Computation and Engineering directorates who bring their expertise to the Laboratory's Security and Protection Program. The Argus security kernel (the real-time system) is written in Ada and deployed on legacy DEC/Compaq Alpha-processor-based hardware using the Virtual Memory System (VMS) operating system. In recent years, a substantial portion of the system (primarily user interfaces) was migrated to a more modern Windows platform.

PROGRESS IN 2007

In the summer of 2007, the Argus team released a major upgrade named Baseline 25, which responded to more than 60 software change requests. A major renovation of the access control system architecture also began in 2007 and is scheduled for release in June 2008. This effort improves the user interface by introducing a Windows-based

graphical user interface. Large portions of the host database software architecture are being redesigned to upgrade data management from the aging Ingres database to a contemporary Oracle database. These upgrades modernize Argus's underlying technology without significantly affecting customers.

The Argus Field Processor (AFP) replacement and Host Migration projects also made significant progress in 2007. Scheduled for release in February 2008, the AFP replacement is a multiyear project to replace obsolete hardware, particularly the main processing boards, and update the AFP software to support the new hardware. The Host Migration project will move the control systems from Alphas to Hewlett-Packard Itaniums and will replace the VMS operating system with a Linux/Unix operating system. This three-year multimillion-dollar project is essential to ensuring the future viability of the Argus system.

SIGNIFICANCE

Argus security helps ensure that the Laboratory's Top Secret documents, materials, and facilities are thoroughly protected, intruders can be detected in real time, and intrusions and emergencies get immediate response from police and investigative personnel. With Argus, LLNL and other DOE sites are provided with 24/7 maximum security. The continual upgrades and enhancements that are made to Argus will help protect the nation's significant nuclear assets and information.



4.08 — ALE3D Sets Its Sights on the Future

CONTACT INFORMATION Rob Neely, rneely@llnl.gov



OVERVIEW

Since the ALE3D project began at LLNL almost 20 years ago, it has undergone continual change to support new architectures, languages, user requirements, and advanced physics. The multidisciplinary ALE3D team, comprising employees from the CMELS, and Engineering directorates, and the Weapons and Complex Integration Principal Directorate provides unique simulation capabilities to LLNL scientists and to NNSA and DoD partners. In addition to the usual litany of feature requests, bug fixes, performance, and fidelity improvements, 2007 saw the continuation of a “rebuild in flight” philosophy aimed at ensuring that ALE3D will be a viable code development platform for many years to come, providing users with a solid and robust code base with which to meet programmatic goals.

PROGRESS IN 2007


Among the notable changes in 2007 is the official release of a two-dimensional (2D) capability in ALE3D. Because of symmetries in the geometry and boundary conditions as a problem evolves over time, some simulations can be performed effectively in 2D. These simulations can be calculated significantly faster and/or with higher resolution than their full three-dimensional (3D) counterparts. Hence, calculations that once took days to complete can be accurately approximated in hours or even minutes. The net effect is a much faster workflow for end users,

who often use 2D to work out details on a simplified mesh before advancing, with only minor changes to the input, to the full 3D model.

The ALE3D team reached a milestone with the adoption of Vista, a standardized data model used in the core architecture of ALE3D. Developed over the past several years as a standalone library and incrementally introduced over time, Vista is now used as the basis for most major data structures in ALE3D. Vista provides a unified data storage and access model implemented as base classes for what were once ad hoc data structures. This model allows natural abstractions for managing and maintaining complex mesh and field data structures. The adoption of Vista immediately began to pay off by setting the stage for future developments, such as adaptive mesh refinement, dynamic load balancing, and easier integration of new physics capabilities (such as a new magnetohydrodynamics package).

ALE3D also continues to grow its user base. Based in the Weapons and Complex Integration Principal Directorate, it has long been a workhorse code for many LLNL directorates. Through the Laboratory’s Joint Munitions Program, ALE3D is increasingly being adopted by DoD partners and contractors. Once or twice a year, an official distribution compact disc is created and sent to more than 40 sites, each of which supports a number of internal users.

This 2D axisymmetric calculation (revolved to show full geometry) runs hundreds of times faster than similarly resolved quarter-symmetry 3D calculations.



SIGNIFICANCE

As multiphysics, multidimensional codes become more complex—supporting coupled physics models, new architectures, and user demands for robust software—solid computer science is essential. Working as part of an integrated multidisciplinary team, ALE3D’s computer scientists contribute directly to the algorithm and model development of the code. They also determine which infrastructure needs are necessary to ensure that the code base will survive decades of development with continuous and incremental improvements.



4.09 — A Context-Aware Nuclear Evaluation System Increases Security at Ports of Entry

CONTACT INFORMATION

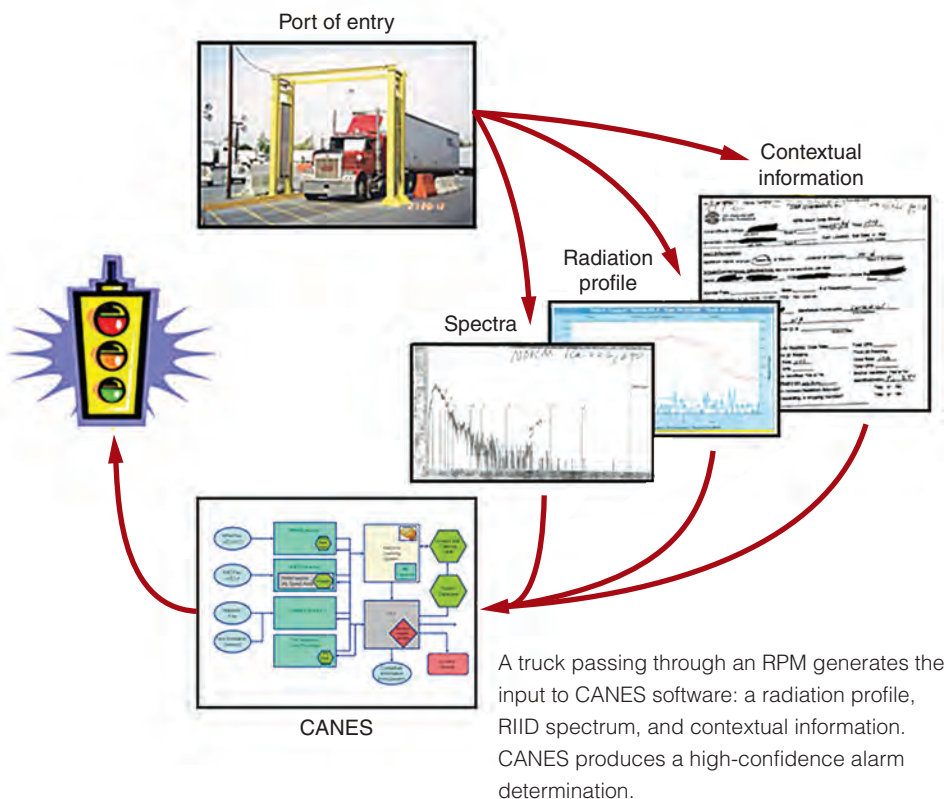
Kelley Herndon Ford, herndonford1@llnl.gov
Simon Labov (Physical Sciences), slabov@llnl.gov



OVERVIEW

The ideal radiation detector system for homeland security would provide an unambiguous signal indicating the presence of a nuclear weapon, special nuclear materials, or a radiological weapon. In some cases, differentiation between signals from benign sources (such as naturally occurring radioactive material) and those from contraband agents cannot be achieved solely through spectroscopic classification and identification. Although

considerable effort has been devoted to improving detector performance, Laboratory scientists are researching ways to better use existing hardware by incorporating contextual information (such as source distance and the type of conveyance) into the analysis of detector data. The team is developing the Context-Aware Nuclear Evaluation System (CANES), which combines multiple data sources and uses machine-learning algorithms to classify new incidents.



PROGRESS IN 2007

LLNL completed a prototype system that combines radiation measurements with contextual information to extract relevant features and uses machine learning to classify radiation alarms. This system was trained on semi-synthetic incidents based on historical data. As part of a user study, team members visited ports of entry to ensure that CANES can be successfully integrated into field operations.

CANES combines information from several types of radiation detectors from different manufacturers, including radiation portal monitors (RPMs), handheld radioactive isotope identification devices (RIIDs), and personal radiation dosimeters. The feature extractors use a physical understanding of the measurements provided by the detector and present results that can be easily understood by nuclear analysts. For example, the RPM extractor provides the intensity of the peak to background

and the shape of the profile. The RIID extractor uses multiple linear regression analysis on the spectra to determine which radionuclides are present, their activities, and the amount of shielding between the source and the detector.

The machine-learning library was developed by Professor Artur Dubrawski's group in the Auton Lab at Carnegie Mellon University. They investigated several methods, including those based on Kernel Density Estimation, Bayesian Belief Networks, and Nearest Neighbor approaches. Machine learning requires a sufficient number of prior incidents (both threat and non-threat data) to adequately train the model. However, obtaining large historical datasets is difficult since radiation screening has only become widespread in recent years. The team developed tools that use domain knowledge to vary historical incidents and generated 6,000 semi-synthetic cases they then used to demonstrate the system. They continue to gather additional datasets to further evaluate the system.

SIGNIFICANCE

CANES will improve threat/non-threat discrimination of radiation alarms, increase confidence in alarm adjudications, lower false-alarm rates, and enhance radiation detection sensitivity at the nation's ports of entry. In addition, the use of machine learning will allow scientists at LLNL and Carnegie Mellon University to develop algorithms for anomaly detection and consistency checking. CANES will reduce operating costs, duplicative data entry, and the time required to adjudicate alarms for both field agents and reachback analysts.

return
to TOC

5

Computation's Institutional Contributions Make a Difference across the Laboratory

In addition to the missions in high-performance computing, programmatic support, fundamental research, and applications development described earlier in this report, the Computation Directorate plays a critical role in meeting the information technology (IT) needs of Lawrence Livermore National Laboratory (LLNL) as a whole. Computation continues to design, direct, enhance, and deliver capabilities others use to accomplish institutional and programmatic needs. This is a complex mission because of the widely disparate needs of the many clients and the broad diversity of activities these efforts encompass. Institutional requirements come from the Director, the Office of the Chief Information Officer, LLNL's Safeguards and Security and Cyber Security programs, the individual programs, and Computation itself. Many efforts arise from a combination of sources. The diversity, scope, and impact of these institutional support activities can be demonstrated by considering just a few examples. Three large-scale efforts are presented in this section: institutional computing, IT systems used across the Laboratory, and computer and network security and accreditation. Several smaller and sometimes more program-specific efforts are also in the sampling, including server virtualization, Advanced Simulation

5.00 — Computation's Institutional Contributions Make a Difference across the Laboratory

CONTACT INFORMATION Ronald Schmucker, schmucker1@llnl.gov



and Computing (ASC) Verification and Validation (V&V), Computation's internal technology development (techbase) efforts, and Computation's efforts in recruiting and academic outreach.

Computation provides 81 trillion floating-point operations per second (TF) of unclassified computing capacity that is used by approximately 2,000 scientists. While some of these scientists are at LLNL, many affiliates and collaborators around the world connect remotely to perform calculations. Numerous LLNL science efforts are advanced with the knowledge gained using these computational and visualization resources.

Providing a large and diverse laboratory with IT systems has always been a challenge for LLNL. The ever-present reality of rapidly changing technology creates a constant dilemma in balancing productivity with the need for seamless and integrated solutions and the strict security requirements of a national laboratory. In addition, LLNL is focusing on making significant efficiency gains in the IT arena, requiring Computation to substantially change IT practices and concentrate on cost-effective delivery while still meeting customer expectations and program needs. Computation is implementing the Information Technology Infrastructure Library (ITIL) standard for defining and delivering IT services. ITIL provides a framework for baselining, service-level agreements, and change management. With service-level agreements, Computation plans to use the ITIL framework to organize and prioritize customer needs, gaining efficiency in IT through more structured central management than in the past.

Computation has devoted significant effort to and achieved considerable success in providing computer security at the Laboratory. For example, by scanning LLNL e-mail and initiating egress port blocking, the directorate better controls the information that leaves the Laboratory, effectively eliminating the problem of spyware sending information off site. Egress port blocking, in conjunction with blocking unsafe Web sites and scanning incoming Web material for malicious content, has closed off commonly used vectors for much of the malicious software (malware) delivery. Malware

delivery techniques continually evolve; recently, they have gone from being entirely e-mail-centric to delivery through Web content and removable media. This evolution necessitates the continual development of increasingly sophisticated preventive measures. At LLNL, the scanning of Web mail to close that malware delivery path is scheduled to begin in early 2008. Despite evidence of a significant number of intrusion attempts, stronger authentication and other security efforts resulted in LLNL having no major computer disruptions and no known intrusions in 2007.


Computation is also responsible for numerous smaller efforts in the institutional computing arena. One of them is the techbase effort within the directorate. Computation dedicates both personnel and funding to continually improve core computer science competencies and capabilities, responding to current and anticipating future needs of the programs. In 2007, the directorate increased techbase spending over prior years and conducted projects to enhance core capabilities in the areas of simulation, data science, and biocomputing.

Computation's contributions to simulation date back many years. For the last decade, driven largely by ASC, LLNL programs, and the Office of Science's Scientific Discovery through Advanced Computing Program, this area has been the major focus of LLNL's computational science efforts. The development of the massively parallel supercomputers that enabled great leaps in simulation capabilities also caused the generation of datasets of extraordinary size and complexity. These data, together with the phenomenally large and intricate datasets generated in national security and bioscience applications, have given rise to extremely difficult issues in data science involving the size of datasets, the rate at which data are acquired and analyzed, and the handling of disparate data sources. Biocomputing is also an important and growing niche, and Computation invested in several projects to grow LLNL's capability in this rich field.

The articles in this section do not give a comprehensive view—they merely scratch the surface of how Computation supports LLNL's institutional activities and illustrate the increasing importance of efficiency in doing that work.

5.01 — ASC V&V Best Practices Yield Better Software at Lower Cost

CONTACT INFORMATION Greg Pope, pope12@llnl.gov



OVERVIEW

Ongoing ASC V&V Software Quality Engineering (SQE) Best Practice Forums, assessments, and other software process improvement initiatives resulted in the 2007 roll-out of important new tools and techniques for SQE on ASC codes. Specifically, recent efforts created a new automated static analyzer tool, increased SQE oversight for feeder and library codes, and centralized SQE project information and compliance documentation. The collaborative efforts of the SQE group and code teams have produced measurable improvement in code quality and, by removing certain types of defects early, may have avoided as much as \$11M in future defect detection and correction costs.

PROGRESS IN 2007

In recent years, the ASC software quality engineers have collaborated closely with the ASC code teams to find and share best practices across various projects. These efforts dramatically enhance the quality of scientific research codes while ensuring compliance with Department of Energy (DOE), National Nuclear Security Administration (NNSA), institutional, and ASC standards and orders. More than 35 Best Practices Forums have been held since September 2004; in these meetings, code team members share effective practices for developing software and warn of ineffective or counterproductive approaches.

One key contribution in 2007 is the addition to the ASC V&V program of



Klocworks K-7, an automated static analyzer. This new tool, which was procured from an outside vendor, was first deployed in November 2006 and has since analyzed more than two million lines of C and C++ code, helping the ASC code teams identify and correct more than 3,000 code defects. Examples of the most commonly found defects are variables that are not initialized prior to first usage, memory that is allocated but subsequently not deallocated (memory leak), overflow of fields by storing array elements or strings beyond the allocated field length, null pointer dereferencing, values that are never used after assignment, dead code, and loss of precision. Many of these defects would have been difficult to find using peer reviews or testing techniques.

In the past year, SQE oversight increased, focusing on the smaller feeder and library codes, which provide important features and data to the larger simulation codes. To help these smaller teams achieve higher code quality ratings,

A physics code gets “Klocworked.” LLNL Software Quality Engineers Ellen Hill, Juan Hernandez, and Bill Oliver (seated) confer with LLNL Physicist Jeff Grandy, who is pointing out a finding in the Klocworks report for an ASC physics code.

the ASC SQE team added a number of configuration management tools, issue trackers, automated test suites, and build/make scripts. In addition, ASC adopted the SourceForge tool to better track the 35 ASC code projects.

SIGNIFICANCE

It costs approximately \$100 to correct each defect found using the Klocworks K-7 static analyzer tool. Comparatively, it costs approximately \$4K to deal with defects users find later, according to industry studies. This difference implies that in 2007 as much as \$11.7M in potential future costs may have been avoided by using the static analyzer. The capability to analyze defect trends from the static analyzer's defect reports also enables the SQE group and code teams to identify the root causes of defects and prioritize process improvements to prevent similar defects in the future. Attaining improved software quality, while avoiding significant future costs, was an important achievement in 2007.

return
to TOC

5.02 — Institutional Computing Advances Grand Challenge Science

OVERVIEW

Bringing its extraordinary computing power to bear on a broad collection of important science questions, Institutional Computing at LLNL allocated 83.7 million hours of central processing units for 17 “Grand Challenge” research projects. The breadth of the program is remarkable, covering topics from astrophysics to chemistry, materials science, bioscience, earthquakes, and climate change. The Peloton procurement in 2006 significantly increased both capability and capacity cycles to further enhance institutional research. Grand Challenge recipients were also allocated time on Atlas, the new 44-TF machine, and on the 22-TF Thunder machine. These systems are dedicated to unclassified research through the Laboratory’s Multiprogrammatic and Institutional Computing Program.

Making high-performance computing resources available to compelling science projects is important to DOE, NNSA, and many LLNL missions. Cherry Murray, LLNL’s principal associate director for Science and Technology, says, “Our institutional program aligns with national security needs and serves as a complement to the DOE Office of Science effort to allocate time on high-performance computing systems to research projects of vital national interest in energy, environment, bioscience, and climate change.”


PROGRESS IN 2007

In February 2007, Murray worked with members of the Laboratory Strategic Program Board and the Laboratory Science and Technology Office to select the current set of Grand Challenge projects. To be considered, proposals had to “address a grand-challenge-scale, mission-related problem that promises unprecedented discoveries in a particular scientific and/or engineering field of research, and if successful, will result in high-level recognition by the scientific community at large.”

In one project, scientists from several institutions are conducting the most detailed multidecade global climate simulations ever performed; this work is part of the larger effort to understand how climate changes over time (see Section 3.05). The team, comprising scientists from LLNL, the Scripps Institution of Oceanography, the National Center for Atmospheric Research, the University of Michigan, and Los Alamos and Oak Ridge national laboratories, is performing climate simulations of unprecedented detail and fidelity and then making them available to the climate research community. With this project, the team has the potential to dramatically advance scientists’ understanding of climate and how it may be affected by anthropogenic (human-induced) factors.

Another project aims to create the capability to predict the backscatter of laser light as it hits ignition targets in experiments at the National Ignition Facility (NIF). This calculation is inherently difficult because backscatter is a nonlinear, multidimensional process that spans from the macroscopic length and time scales of hydrodynamics to the microscopic scales of electrons. In simulating beam propagation

and backscatter, the team uses radiation-hydrodynamics simulations to generate plasma profiles at peak laser power and inputs these into pF3D, LLNL’s three-dimensional (3D) beam propagation code. The pF3D code operates on a spatial scale of micrometers (10^{-6} meters) and a time scale of picoseconds (10^{-12} seconds). Thus, a few cubic millimeters in space corresponds to billions of calculation zones, which in turn



In this screen shot from a backscatter simulation, the laser beam has propagated up to the higher-density region near the hohlraum wall.




CONTACT INFORMATION

Brian Carnes, carnes1@llnl.gov
 Web site: computing.llnl.gov/mic/

translates into tens of thousands of processor-hours per 3D simulation. With access to the Grand Challenge time and equipment, the scientists can now include the effects of transverse gradients in the plasma profiles and simulate scattering in multiple materials.

These types of calculations, which represent a major advance in modeling capabilities, would not have been possible without the Grand Challenge allocation—an oft-repeated

story from program recipients. Livermore's Grand Challenge Program is the logical outgrowth of a long-term effort to support unclassified computational science. During the past 10 years, the high-performance computing resources dedicated to the Laboratory's unclassified institutional research have increased more than a thousand-fold, from 72 billion floating-point operations per second (or 0.072 TF) in 1997 to 81 TF today.



This screen shot shows the backscatter-stimulated Raman scattering (green) growing and spreading near the hohlraum wall and propagating back toward the laser entrance hole. Stimulated Brillouin scatter (blue) is beginning to grow and slightly spread near the capsule.

SIGNIFICANCE

High-performance computing is an integral part of the scientific process and can serve to advance science in areas of vital interest to the country and the global community, including energy, bioscience and medicine, and climate change. The new high-performance computing resources Computation has put in place for unclassified research leverage the directorate's supercomputing expertise and infrastructure for stockpile stewardship. The Grand Challenge projects will, in turn, benefit the Laboratory's national security missions. Institutional computing is an essential component of LLNL's science and technology investment strategy and has been instrumental in helping the Laboratory generate headline scientific and technical accomplishments that benefit the nation and the global community.

The 17 Grand Challenge research projects awarded time on institutional computing resources in 2007 were:

- Ultra-High-Resolution Coupled Climate Simulations
- Toward a Predictive Capability for Laser Backscatter in NIF Ignition Targets
- Dislocation Patterns: Fractals or Cells?
- Study of Ramp Compression of Metals Using Molecular Dynamic Simulations
- Supernova Grand Challenges
- Predicting Properties of Plutonium with Dynamical Mean Field Theory
- Toward Petascale Atomistic Simulations with Quantum-Level Accuracy
- From Quantum Chromodynamics to Nuclei to Stars
- Atomic Simulation of Shockless Compression
- Accretion Disks around Rapidly Rotating Black Holes
- Long-Time-Scale Shock Dynamics of Astrophysical Ices
- Understanding Multicomponent Heterogeneous Processes in Atmospheric Chemistry
- Simulation of Membrane Protein and Nanodisc Particles
- Gyrokinetic Simulation of Turbulence and Transport in Tokamak Plasmas
- Advanced Explosion and Earthquake Simulations
- Full-Scale Simulations of Reduced Mass Targets and Wakefield Accelerators
- Quantum Simulation of Nanomaterials under Pressure



5.03 — Academic Interactions Produce Research and Recruits

CONTACT INFORMATION


David Keyes, dekeyes@llnl.gov
 James McGraw, jmcgraw@llnl.gov
 Web site: computation.llnl.gov/iscr/



OVERVIEW

Fulfilling LLNL's missions requires dramatic progress across a broad spectrum of computational science research topics. Many technical advances can only be made through extensive collaborations with academia. The talent and unfettered enthusiasm of academic researchers pair effectively with LLNL's focused, mission-oriented strategy and, not insignificantly, with the Laboratory's access to the world's most powerful computers.

Through its Institute for Scientific Computing Research (ISCR), the Computation Directorate builds and maintains bridges to the academic community with a program of managed research, collaborative subcontracts, faculty visits, student internships, workshops, and seminar series. ISCR hosts both short- and long-term visits, with activity peaking during the summer months.



Annette Lettsome (Michigan State University) describes her summer project on enabling the simultaneous use of multiple graphics windows to enhance the visualization capabilities for a physics simulation code called Ares.

PROGRESS IN 2007

During 2007, ISCR supervised a joint effort between scientists from LLNL and UC San Diego. Called the Livermore–UC San Diego Scientific Data (LUSciD) study, this effort is identifying and applying the mechanisms needed to conduct collaborative research involving massive datasets for extremely large-scale scientific simulation. LUSciD projects focus on three topics: regional climate modeling, astrophysics, and scientific data management. Results from a 2007 LUSciD analysis, for example, showed a human-induced signal in the histories of snowpack, river flow, and temperature in the western United States.

ISCR also coordinates the University Education Partnerships Program, which in 2007 continued to fund educational projects at UC campuses in Berkeley, Los Angeles, and San Diego and at San Diego State University. On these four projects, a faculty principal investigator and a graduate student work closely with an LLNL collaborator.

Seventy students, about half of whom were doctoral candidates, participated in the 2007 Summer Scholar Program.

Typically in residence for 12 weeks, these students work with an LLNL mentor and other Laboratory scientists conducting research or developing software in support of Computation's missions. Projects support a range of Laboratory research, from NIF to the biological sciences, physics, electrical engineering, chemistry, aeronautics and astronautics, mechanical engineering, information science and systems, and, of course, computer science and applied mathematics. Although the students' primary activity is research, they are encouraged to attend as many seminars as possible to learn more about the Laboratory's unique environment.

The Institute for Terascale Simulation Distinguished Lecturer Program supports lectures and technical discussions with international leaders in the foundations of simulation. In 2007, ISCR hosted J. Tinsley Oden, a professor from the University of Texas; Cleve Moler of MathWorks; and Robert Rosner, director of Argonne National Laboratory. In addition to the summer students, ISCR hosted 189 short-term visits for Computation, including 55 seminars.

SIGNIFICANCE

By managing joint research, identifying and hosting researchers from the academic community, and operating an active internship program, ISCR fosters computational science collaborations for the Computation Directorate. Through this work, ISCR provides long-term academic research projects that address LLNL's mission priorities and help cultivate the Laboratory's future workforce.

5.04 — Server Virtualization Reduces Costs in Weapons and Complex Integration Organization

CONTACT INFORMATION Scott Taylor, taylor69@llnl.gov



OVERVIEW

The Weapons and Complex Integration (WCI) Computer Support Unit within Computation recently completed a project to consolidate multiple servers onto a few physical machines using computer virtualization technology. Computer virtualization allows multiple operating systems to run simultaneously on a single physical machine, which reduces hardware while increasing both hardware utilization and reliability. Each separate virtual machine operates as though it has its own processors and memory and is treated as a distinct system. At the same time, it can be easily moved to different physical hardware.

The project was launched using ESX Server 2.1 from VMware and had three main goals: reduce the number of physical computers (to save on hardware costs), reduce the amount of time required to deploy and manage new systems (to save on support costs), and simplify hardware upgrade procedures. Initial cost estimates indicated that the project would be justified if five virtual systems per physical system could be deployed.


PROGRESS IN 2007

The pilot project began in mid-2006 when four physical servers were purchased and configured to run ESX Server 2.1. Approximately 20 servers that were not considered mission-critical were migrated from individual physical hardware to virtual machines running on the four new physical servers. The initial deployment was successful, and the decision was made to expand the project.

In 2007, the pilot became a production project and passed its first true agility test in the WCI Principal Directorate. Computation's WCI Computer Support Unit deployed a core server and several participating servers and workstations as virtual machines. In just eight hours and with no procurement cost, 17 systems were successfully deployed on a single physical server. Later in 2007, some of the initial testing systems were upgraded with faster hardware and VMware ESX 2.5. These upgrades included installing the updated virtualization software on the new hardware and moving existing virtual machines to the new physical equipment. The Computer Support Unit accomplished the move in a few minutes, with little or no down time for the virtual machines, by using the live-migration capabilities of VMware. Several more systems, including some mission-critical systems, were either migrated to or deployed as virtual machines in 2007.

The additional deployment of virtual machines in WCI brings the total to 74 virtual machines that run on nine physical servers. WCI's experience, based on the hardware currently deployed, has shown that 10 to 15 virtual servers can run on

each physical machine, greatly exceeding the initial goal of five virtual machines per physical server. The time required to deploy new systems has been reduced from weeks to hours, and power and cooling requirements have also decreased significantly.



Server virtualization cost comparison in the WCI Principal Directorate.

SIGNIFICANCE

The experience in WCI demonstrates that the benefits of virtualization accrue even in a highly specialized scientific software development and testing environment (including file servers, database servers, Web/application servers, Windows terminal servers, compile/development servers, and small-scale compute servers). The results make a strong case for aggressive, wide-scale virtualization of programmatic and institutional servers at LLNL.



5.05 — Cyber Security Protects Against Evolving Threats to LLNL's Computing Assets

CONTACT INFORMATION Ed Matsche, matsche1@llnl.gov



OVERVIEW

Personnel in the Information Technology Protection Division (ITPD) support the Laboratory's Cyber Security Program by overseeing cyber security aspects of LLNL's computing assets, ensuring compliance with NNSA's classified and unclassified computer and telecommunications security requirements, and managing LLNL's institutional cyber security technology infrastructure.

PROGRESS IN 2007


In 2007, ITPD staff expanded LLNL's suite of network protection software to cover the increasing threats associated with Web access and Web site vulnerabilities. An outbound Web proxy service (Blue Coat) and intrusion prevention system (Tipping Point) protect LLNL users from accessing potentially unsafe Web sites that automatically exploit vulnerabilities and execute arbitrary code on a system without user knowledge. The Web proxy also protects against malware by scanning downloaded Web pages for malicious code and blocking infected sites. To counter the increasing cyber security threat posed by Web applications, ITPD also began scanning key Web servers for application-level vulnerabilities, enabling developers to counter new attack vectors.

LLNL has long used the remote access methods Institutional Port Allow and Web Proxy Service, which opened ports into the internal network from a specific

Internet Protocol address. These methods were replaced in 2007 with a more secure virtual private network technology using Secure Socket Layer (SSL). This solution requires fewer institutional support resources while enhancing security. Employees on business travel who need remote access can now use any Web-based browser that supports Java and SSL. Information downloaded during the session is encrypted and automatically removed from the computer after the user logs out. In addition, this new technology enabled LLNL's Highly Enriched Uranium Transparency Program to retire a firewall policy exception and better protect its confidential data.

LLNL's Entrust Support Team executed "Government 2 Business" contracts with several outside organizations, including the Livermore Police Department and UC Berkeley. Under the contracts, LLNL can authorize individuals in these organizations to obtain Entrust identities and certificates from DOE Headquarters, enabling them to securely exchange unclassified controlled information via encrypted e-mail with LLNL (or other DOE and Department of Defense) collaborators.

ITPD personnel oversaw the implementation of port restrictions on LLNL's classified computers by creating documents that define LLNL standards for the required controls and validating the implementations through internal and Livermore Site Office spot-checks.



Incident Management Team members Willem Verschuur, Erwin Lopez, and Matthew Myrick display a real-time dashboard of attacks intercepted by the Blue Coat Web proxy and Tipping Point intrusion prevention system managed by ITPD staff.

SIGNIFICANCE

LLNL experienced numerous attacks against its computers in 2007 but very few successful intrusions. The Web proxy prevented approximately 750 attempts to download malicious code, 98,000 attempts to access Web sites that contain spyware, and more than six million attempts to visit Web sites that host malicious content. ITPD's efforts to minimize the paperwork and delay in accrediting systems helped reduce the number of classified information security plans from 68 to 53 and resulted in one unclassified master security plan that restored local authority to approve individual systems and applications. ITPD staff collaborated with other areas of the Computation Directorate to test and certify enabling technologies such as virtual local area networks, network access control, and role-based access control, laying the foundation for the future Blue Network, which will address the Laboratory's unique security requirements for foreign nationals.

return
to TOC

5.06 — Production Begins on Institutional Wireless Network

CONTACT INFORMATION

Robin Goldstone, goldstone1@llnl.gov



OVERVIEW

The LLNL Wireless Network was approved for production deployment in early 2007 and is now being managed as an accredited institutional service, available for use by all LLNL employees. The service is managed by the Backbone Network Team (BNT) within Computation's Integrated Computing and Communications Division. BNT network engineers have designed the LLNL Wireless Network to provide the highest degree of usability and reliability while also meeting the stringent security requirements for NNSA accreditation. BNT, in partnership with the Engineering Directorate's Engineering Manufacturing and Services Section, works closely with customers to design and deploy wireless networking infrastructure in various locations at the Laboratory.

PROGRESS IN 2007

In 2007, the LLNL Wireless Network progressed from a pilot project to a supported production service, requiring technical enhancements, approved security processes, and a defined operational support model. Technical enhancements included a transition from standalone authentication to the use of LLNL's Active Directory service, providing the wireless controller infrastructure with sufficient redundancy that component failure does not cause system failure, and an upgrade of the wireless firewall from handling 100 million to 1 billion bits per second. Streamlined processes allow user account requests and access point approvals to be managed more smoothly, and

coordination with the LLNL Cyber Security Program for wireless intrusion detection has also improved the service. A new Wireless Network Information Portal disseminates information about the service, and the 4-HELP institutional help desk provides support and tracks issues.

There are presently 290 registered users on the LLNL Wireless Network. Access points are installed in 20 buildings. Initial pilot and early production efforts provided wireless network coverage in common areas including LLNL's West and Central cafeterias, the library, and auditoriums. Ongoing efforts are focused on providing network coverage in conference rooms in property-protected buildings throughout the Laboratory as well as in programmatic facilities. NIF has been particularly aggressive in deploying wireless network

coverage throughout its facilities to streamline meetings and increase productivity.

A number of assessments conducted in 2007 explored the feasibility of supporting additional wireless networking capabilities. These assessments included a successful demonstration of visitor wireless access, a model for leveraging existing wireless network infrastructure to provide a wired captive portal solution in conference rooms and other locations where wireless devices are not permitted, and an investigation into the security risks and mitigations of allowing embedded wireless adapters to be used in LLNL-owned laptops. These assessments will pave the way for enhancements to the LLNL Wireless Network service in 2008 and beyond.

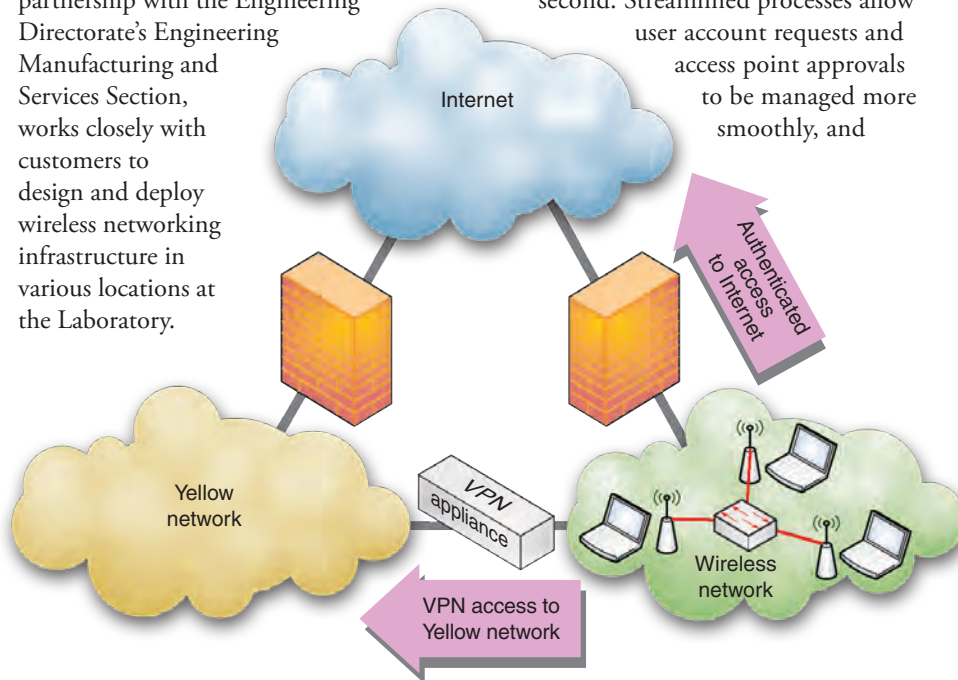


Diagram of the LLNL Wireless Network.

SIGNIFICANCE

BNT made strides in 2007 to provide the LLNL user community with a true mobility solution that uses a single authentication domain and dynamic Internet Protocol addressing for access across the LLNL site. Additionally, next-generation wireless standards offer the promise of bandwidth on par or exceeding that of the existing wired infrastructure in many LLNL buildings. LLNL's wireless networking initiatives offer a flexible, productive, and cost-effective network connectivity solution for a variety of LLNL users.



5.07 — Process Automation Enables Centralized Desktop Configuration Management

CONTACT INFORMATION David Frye, frye3@llnl.gov



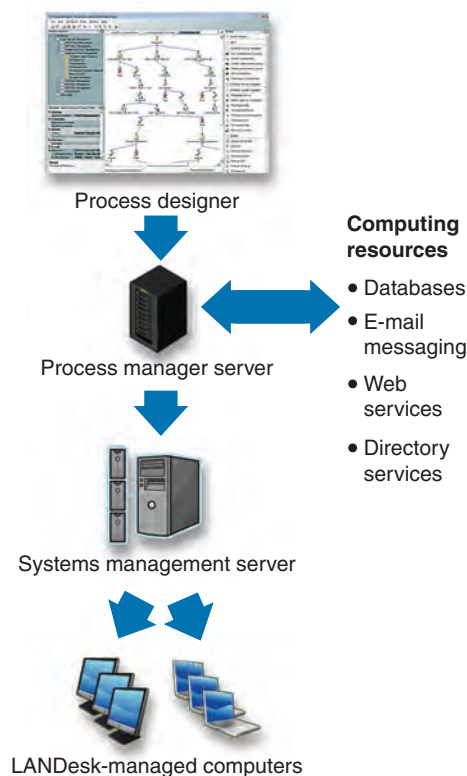
OVERVIEW

In the last several years, LLNL has made substantial investments to improve its unclassified systems management capabilities. The LANDesk Enterprise Management Suite, purchased in 2006, is now deployed to more than 12,000 Windows, Macintosh, and Linux desktop computers and servers on the Laboratory's unclassified network, providing automated control of LLNL's desktop and server infrastructure.

The LANDesk Enterprise Management Suite includes tools for hardware and software inventory management, security update management (patching), automated software distribution, remote desktop administration, and reporting. LLNL now seeks to standardize and automate various IT processes by implementing the LANDesk Process Manager. The Process Manager uses the LANDesk Enterprise Management Suite to provide process-driven automation of the entire system.

PROGRESS IN 2007

In 2007, LANDesk experts in Computation's System Management Solutions Group (SMSG) worked with the systems administrators for each directorate to define patch management processes, develop workflows, and extend the capabilities of the Process Manager



The LANDesk Process Manager orchestrates workflow execution and configures LANDesk-managed computers through the Laboratory's existing infrastructure.

to meet the needs of the Laboratory and its directorates. SMSG developed a hierarchical workflow system consisting of institutional and directorate-specific processes for patch management.

The workflow components provide centralized testing and approval capabilities. When a security upgrade is available, patch content is automatically downloaded and then tested twice—first for compatibility with standard desktop configurations and then against a representative set of computers across the institution. If both tests are successful, the patch content is submitted for general deployment approval.


SMSG also collaborated with systems administrators around the Laboratory to

define patch processes for each directorate. Because patch deployment can affect work activities, the group wanted to ensure that the automation process for scheduling updates, notifying users, and executing the deployment is consistent with the manual processes each directorate has in place. Conversely, directorate processes leverage the rigorous institutional testing processes previously described.

In 2008, SMSG will complete the initial definition and production deployment for the patch management process. In addition, the group will develop workflow-based solutions for other capabilities, such as managing upgrades to desktop software and streamlining the Laboratory's certification and accreditation efforts.

SIGNIFICANCE

The threats to LLNL's computer security from outside hackers and determined adversaries continue to rise as the world becomes increasingly connected, mobile, and technologically sophisticated. Automated patch management through workflow promises to improve the Laboratory's ability to detect threats and automatically initiate the appropriate countermeasures, with little or no human intervention, in a manner that meets the needs of the institution's diverse organizations. Together, the LANDesk Enterprise Management Suite and Process Manager provide a platform for efficiently implementing this new approach, which will dramatically improve LLNL's computer security and free systems administrators to perform more complex, higher-impact work.



6.01 — Computation Directorate Industrial Collaborators

Company	Topic	LLNL contact
Appro	Scalable capacity clusters	Mark Seager
Barcelona Supercomputer Center	SLURM resource management software	Morris Jette
Bull	SLURM resource management software	Morris Jette
CEI–EnSight	Visualization package	Eric Brugger
Cluster Resources, Inc.	MOAB workload manager	Don Lipari
Conoco–Phillips	Oil exploration	Shawn Larsen
Cray	Scalable systems, open-source strategy	Mark Seager
DataDirect Networks	Raid 6 R&D for input/output systems	Brent Gorda
Exegy	Hardware regular expression matcher	John May
FusionIO	Architectures for data-intensive computing	Maya Gokhale
Hewlett–Packard	SLURM resource management software	Morris Jette
IBM	High-performance storage system	Mark Gary
IBM	Scalable systems, multiple areas	Mark Seager and others
IBM	Seismic modeling	Shawn Larsen
IQuum	DNA signature target design	Tom Slezak
Limit Point Systems	LibSheaf	Bill Arrighi
Microfluidic Systems	Pathogen signatures	Tom Slezak
Netezza	Parallel database hardware and software	Scott Kohn
OpenMP Consortium	Shared-memory programming models	Bronis de Supinski
OpenWorks	Valgrind memory tool development	John Gyllenhaal
Panasas/FAST	Architectures for data-intensive computing	John May
ParaTools, Inc.	Performance tools for high-performance computing systems	Chris Chembreau
R3vis Corporation	Chromium render server	Eric Brugger
Red Hat	Operating systems	Mark Seager
Schlumberger	Synthetic seismic dataset generation	Shawn Larsen
Sicortex	SLURM resource management software	Morris Jette
Society of Exploration Geophysicists	Advanced subsurface modeling	Shawn Larsen
Sun Microsystems	Lustre file system deployment	Brent Gorda
TotalView Technologies	Parallel debugger scalability and enhanced memory tools	Dong Ahn and Scott Futral
Tungsten Graphics	Chromium render server	Eric Brugger
URS Corporation	Earthquake hazard estimation	Shawn Larsen
Violin Memory	Architectures for data-intensive computing	Maya Gokhale
Voltaire	OFED: Open-source software stack for InfiniBand networks	Matt Leininger

6.02 — Academic Outreach

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC Berkeley	Chris McKee	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
UC Berkeley	Doug Dreger	Subcontract	Earthquake hazard	IGPP	Shawn Larsen
UC Berkeley	James Demmel	Joint Research	Terascale-optimal partial differential equation solvers	ASCR SciDAC	Rob Falgout
UC Berkeley	Jonathan Wurtele	Subcontract	Theoretical and numerical investigations of Raman backscatter	LLNL Overhead	Richard Berger
UC Davis	Bernd Hamann	Joint Research	Analysis and visualization of scientific data using topology-based methods	LDRD	Valerio Pascucci
UC Davis	Dave Wittman	Joint Research	Large Synoptic Survey Telescope long-term data analysis	LDRD	Don Dossa
UC Davis	Greg Miller	Subcontract	Algorithm development for interface tracking in a compressible medium	ASCR SciDAC	David Trebotich
UC Davis	Kenneth Joy	Joint Research	Real-time three-dimensional viewing of large time-dependent aerial imagery	Sonoma	Mark Duchaineau
UC Davis	Kwan-Liu Ma	Subcontract	Large-graph visualization for homeland security	DOE OCIO	Tony Bartoletti
UC Davis	Kwan-Liu Ma	Joint Research	Large-graph data visualization	LDRD	Valerio Pascucci
UC Davis	Zhendong Su	Subcontract	Software security analysis research	LDRD	Dan Quinlan
UC Los Angeles	Heinz-Otto Kreiss	Consultant	Adaptive methods for partial differential equations	ASCR Base	Lori Diachin
UC Los Angeles	Kendall Houk	Subcontract	Using a combined quantum mechanics/molecular mechanics approach with Car–Parrinello molecular dynamics rare-event methods to study enzymatic reaction mechanisms	LLNL Overhead	James McGraw
UC Los Angeles and Weizman Institute (Israel)	Achiezer Brandt	Consultant	Geometric and algebraic multigrid techniques	ASCR Base	Rob Falgout
UC San Diego	Allan Snavey, Laura Carrington	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
UC San Diego	Dan Reynolds	Joint Research	Implicit solvers and preconditioning techniques for simulations of magnetohydrodynamics and core-collapse supernovae	ASCR Base	Carol Woodward
UC San Diego	David Benson	Subcontract	A predictive model of fragmentation using adaptive mesh refinement and hierarchical material model	LDRD	Alice Koniges
UC San Diego	Joel Conte	Joint Research	Sensitivity analysis for structural engineering and uncertainty quantification	LDRD	Radu Serban
UC San Diego	Randy Bank	Consultant	Numerical solutions of partial differential equations, multilevel iterative methods, and adaptive grid generations	ASC	Rob Falgout
UC San Diego	Reagan Moore, Tim Barnett, Michael Norman, Randy Bank	Subcontract	LLNL–UCSD Scientific Data Management	UC Management Contract	James McGraw
UC San Diego	Scott Baden	Subcontract	Data-driven execution of latency-tolerant algorithms	LLNL Overhead	Dan Quinlan
UC Santa Barbara	Linda Petzold	Joint Research	Numerical methods for sensitivity analysis of differential-algebraic equation systems	ASCR, ASC	Radu Serban

6.02 — Academic Outreach (cont.)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC Santa Cruz	Stan Woosley	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
Atomic Weapons Establishment	Jim Andrews	Joint Research	Debris and shrapnel experiments and modeling	LDRD	Alice Koniges
California Institute of Technology	Daniel Meiron	ASC Alliance	Center for Simulating the Dynamic Response of Materials	ASC	Dick Watson
Cambridge University	Nikos Nikiforkis	Joint Research	Simulation and modeling using Overture	ASCR Base	Bill Henshaw
Carnegie Mellon University	Christos Faloutsos	Subcontract	Group discovery in time-evolving complex networks	LDRD	Tina Eliassi-Rad
Columbia University	David Keyes	Subcontract	Acting Director for the Institute for Scientific Computing Research	LLNL Overhead	James McGraw
Columbia University	David Keyes	Consultant	Algorithms for the solution of partial differential equations on massively parallel computers; general high-performance computing	LLNL Overhead	James McGraw
Columbia University	David Keyes	Joint Research	Terascale-optimal partial differential equation solvers	ASCR SciDAC	Rob Falgout
Commissariat à l'Energie Atomique	Alain Geille	Joint Research	Debris and shrapnel modeling for the National Ignition Facility and the Laser Megajoule	NIF	Alice Koniges
Cornell University	Claire Cardie	Joint Research	Co-reference resolution	DHS	David Buttler
Cornell University	Sally McKee	Subcontract	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Bronis de Supinski
Duke University	Herbert Edelsbrunner	Joint Research	Discrete methods for computing continuous functions	LDRD	Valerio Pascucci
Florida Institute of Technology	Jim Jones	Joint Research	Multigrid approaches for systems of partial differential equations	ASC	Rob Falgout
Georgia Institute of Technology	Jaroslav (Jarek) Rossignac	Subcontract	Advanced prediction for the compression of regular and irregular datasets	LDRD	Peter Lindstrom
Krell Institute	Lucille Kilmer	Subcontract	Department of Energy High-Performance Computer Science Fellowship Program—LLNL portion	ASC	John May
Krell Institute	Tom Brennan	Subcontract	Investigative work on cooperative parallelism	LLNL Overhead	James McGraw
McMaster University	Nedialko Nedialkov	Joint Research	Graph coloring algorithms	ASC	Radu Serban
National Optical Astronomy Observatory	Chris Smith	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LSST Corporation	Don Dossa
National University of Defense Technology	Hongjia Cao	Joint Research	SLURM resource management software	LC	Morris Jette
National Center for Supercomputing Applications	Ray Plant	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LSST Corporation	Don Dossa
New York University	Olof Widlund	Joint Research	Terascale-optimal partial-differential-equation solvers	ASCR SciDAC	Rob Falgout
Ohio State University	Umit Catalyurek	Subcontract	Massive-scale semantic graphs	DHS	James McGraw
Penn State University	James Brannick	Joint Research	Optimal algebraic multigrid interpolation and convergence theory	ASC	Rob Falgout

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Penn State University	Jinchao Xu	Joint Research	Optimal algebraic multigrid interpolation and convergence theory	ASC	Rob Falgout
Penn State University	Ludmil Zikatanov	Subcontract	Optimal algebraic multigrid interpolation and convergence theory	ASC	Rob Falgout
Portland State University	Karen Karavanic	Joint Research	Performance analysis infrastructure for petascale applications	ASC CSSE	John May
Purdue University	David Ebert	Subcontract	Text extraction and data visualization for pet health surveillance	DHS	James McGraw
Purdue University	Ziqiang Cai	Summer Faculty	Multigrid methods for the H(div)-type pseudostress system of the Stokes equations	ASC	Barry Lee
Rensselaer Polytechnic Institute	Don Schwendeman	Subcontract	Development of numerical methods for mathematical models of high-speed reactive and nonreactive flow	ASCR Base	Bill Henshaw
Rensselaer Polytechnic Institute	Ken Jansen	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
Rensselaer Polytechnic Institute	Mark Shephard	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
Rice University	John Mellor-Crummey	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
Rutgers University	Fred Roberts, Paul Kantor	Joint Research	University Affiliate Centers	DHS	James McGraw
Rutgers University	James Abello	Subcontract	Graph Fusion in External Memory	DHS	James McGraw
San Diego State University	Calvin Johnson	Subcontract	Frontier computations in the structure of atomic nuclei	LLNL Overhead	Erich Ormand
Stanford University	Parvis Moin	ASC Alliance	Center for Integrated Turbulence Simulations	ASC	Dick Watson
Stanford University	Pat Hanrahan	Subcontract	General-purpose graphics processing units computing	GAIA	David Bremer, Holger Jones
Stanford University	Stuart Marshall, Mike Huffer	Joint Research	Large Synoptic Survey Telescope camera data output hardware and imaging formats	LDRD	Don Dossa
State University of New York, Stony Brook	Doug Swesty	Joint Research	Solvers for supernova simulation	ASCR SciDAC	Carol Woodward
State University of New York, Stony Brook	Jim Glimm	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
State University of New York, Stony Brook	Michael Zingale	Subcontract	Verification and validation of radiation hydrodynamics for astrophysical applications	LDRD	Louis Howell
State University of New York, Stony Brook	Xiao-Lin Li	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
Technical University of Dresden	Wolfgang Nagel	Joint Research	Parallel performance analysis tools	LC	Chris Chabreau
Texas A&M University	Joseph Pasciak	Subcontract	Algebraic multigrid for problems for h-curl	ASCR Base	Panayot Vassilevski
Texas A&M University	Raytcho Lazarov	Joint Research	Algebraic multigrid for problems for h-curl	ASCR Base	Rob Falgout
University of Arizona	Adam Burrows	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell

6.02 — Academic Outreach (cont.)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Arizona	Tim Axelrod	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LSST Corporation	Don Dossa
University of British Columbia	Carl Olivier-Gooch	Subcontract	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
University of Chicago	Don Lamb	ASC Alliance	Center for Astrophysical Thermonuclear Flashes	ASC	Dick Watson
University of Colorado	John Ruge	Subcontract	Algebraic multigrid, spectral AMGe, FOSPACK, applications to ALE3D problems	ASC	Rob Falgout
University of Colorado	Marian Brezina	Joint Research	Algebraic multigrid, spectral AMGe, FOSPACK, applications to ALE3D problems	ASC	Rob Falgout
University of Colorado	Steve McCormick	Joint Research	Terascale-optimal partial-differential-equation solvers	ASCR SciDAC	Rob Falgout
University of Colorado	Steve McCormick	Subcontract	FOSPACK, AMG, adaptive AMG, compatible relaxation, parallelization, and application codes	ASCR, ASC	Rob Falgout
University of Colorado	Tom Manteuffel	Joint Research	Solution methods for transport problems	ASC	Peter Brown
University of Colorado	Tom Manteuffel	Joint Research	Geometric and algebraic multigrid techniques; FOSLS approach to solving partial differential equations	ASC	Rob Falgout
University of Colorado	Tom Manteuffel	Joint Research	Terascale-optimal partial-differential-equation solvers	ASCR SciDAC	Rob Falgout
University of Houston	Yuriy Fofanov	Joint Research	Genomic algorithms	DHS	Tom Slezak
University of Illinois, Urbana-Champaign	Dan Roth	Subcontract	Populating ontologies: named entities and relations	DHS	James McGraw
University of Illinois, Urbana-Champaign	Dan Roth, David Forsythe, Jiawei Han, Cheng Xiang	Joint Research	University Affiliate Centers	DHS	James McGraw
University of Illinois, Urbana-Champaign	Michael Heath	ASC Alliance	Center for Simulation of Advanced Rockets	ASC	Dick Watson
University of Maryland	Jeff Hollingsworth	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Maryland	Lise Getoor	Subcontract	Role discovery in evolving semantic graphs	LDRD	Tina Eliassi-Rad
University of Michigan	Peter Honeyman	Subcontract	NFSv4 scalable performance, security	ASC CSSE	Steve Louis
University of Minnesota	Douglas Arnold	Participating Institution	Institute for Mathematics and Its Applications	ASC	Van Emden Henson
University of New Mexico	Anil Prinja	Subcontract	Validation and verification for radiation and neutron transport	ASC	Peter Brown
University of New Mexico	Thomas Hagstrom	Joint Research	High-order structure grid methods for wave propagation on complex unbounded domains	ASCR Base	Bill Henshaw
University of North Carolina	Robert Fowler	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Pittsburgh	Janyce Wiebe	Subcontract	Text extraction and data visualization for pet health surveillance	DHS	James McGraw

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Pittsburgh	Janyce Wiebe	Joint Research	University Affiliate Centers	DHS	James McGraw
University of Southern California	Eduard Hovy, Patrick Pantel	Joint Research	University Affiliate Centers	DHS	James McGraw
University of Southern California	Robert Lucas, Mary Hall, Jacqueline Chame	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint Research	Terascale-optimal partial differential equation solvers	ASCR SciDAC	Rob Falgout
University of Tennessee	Jack Dongarra, Shirley Moore	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Texas, Austin	Omar Ghattas	Joint Research	Terascale-optimal partial differential equation solvers	ASCR SciDAC	Rob Falgout
University of Texas, San Antonio	Qing Yi	Subcontract	ROSE compiler support	ASC	Dan Quinlan
University of Utah	Charles Hansen	Consultant	Data exploration, multiresolution scientific data visualization, and algorithm design	ASC CSSE	Mark Duchaineau
University of Utah	Claudio Silva	Joint Research	Studying the topology of point-set surfaces	LDRD	Valerio Pascucci
University of Utah	David Pershing	ASC Alliance	Center for the Simulation of Accidental Fires and Explosions	ASC	Dick Watson
University of Utah	Ellen Riloff	Subcontract	Bioforensics text extraction	DHS	David Buttler
University of Utah	Ellen Riloff	Subcontract	Text extraction and data visualization for pet health surveillance	DHS	James McGraw
University of Utah	Ellen Riloff	Joint Research	University Affiliate Centers	DHS	James McGraw
University of Waterloo	Hans de Sterck	Joint Research	Long-distance interpolation for algebraic multigrid	ASCR SciDAC	Ulrike Yang
University of Wisconsin	Bart Miller	Joint Research	Support for enhanced dyninst testing and initial steps toward open binary editing environment	ASC CSSE	Martin Schulz
University of Wisconsin	Dan Negrut	Joint Research	Implicit integration methods for molecular dynamics	ASCR	Radu Serban
University of Wisconsin	Jason Kraftcheck	Joint Research	Mesquite software development	ASCR Base	Lori Diachin
Worcester Polytechnic Institute	Homer Walker	Joint Research	Development and analysis of iterative methods for large-scale linear and nonlinear algebraic systems	ASC	Carol Woodward



6.03 — Publications

BOOKS AND BOOK CHAPTERS

- Ashby, S. F. and J. M. May (2007), "Multiphysics Simulations and Petascale Computing," in *Petascale Computing: Algorithms and Applications*, D. Bader, Ed. (Chapman & Hall, Boca Raton, FL), pp. 47–66 (UCRL-BOOK-227348).
- Isaacs, T. H., P. L. Chrzanowski, and V. E. Henson (2007), *Making History... Making a Difference: Photographs Celebrating Six Decades of Excellence* (Regents of the University of California), 197 pgs.
- Koniges, A. E., et al. (to appear), "Development Strategies for Modern Predictive Simulation Codes," in *Parallel Computing: Architectures, Algorithms and Applications Volume 15 Advances in Parallel Computing*, C. Bischof, M. Bücker, P. Gibbon, G. Joubert, T. Lippert, B. Mohr, and F. Peters, Eds. (IOS Press, Amsterdam, The Netherlands), pp. 697–704.
- Wissink, A., et al. (to appear), "Adaptive Cartesian Methods for Modeling Airborne Dispersion," in *Advanced Computational Infrastructures for Parallel and Distributed Adaptive Applications*, A. Zomaya, M. Parashar, S. Chandra, X. Li, Eds. (Wiley Book Series on Parallel and Distributed Computing) (UCRL-MI-223156).

JOURNAL PAPERS

- Aagardh, B., et al. (2007), "Ground Motion Modeling of the 1906 San Francisco Earthquake II: Ground Motion Estimates for the 1906 Earthquake and Scenario Events," *B. Seismol. Soc. Am.* (UCRL-228005).
- Amendt, P., et al. (2007), "Assessing the Prospects for Achieving Double-Shell Ignition on the National Ignition Facility Using Vacuum Hohlräume," *Phys. Plasmas* **14** (5).
- Appelo, D. and A. Petersson (2007), "A Stable Finite Difference Method for the Elastic Wave Equation on Complex Geometries with Free Surfaces," *Commun. Comput. Phys.* (UCRL-JRNL-237329).
- Arsenlis, A., et al. (2007), "Enabling Strain Hardening Simulations with Dislocation Dynamics," *Model. Simul. Mater. Sci.* **15** (6), 553–595 (UCRL-JRNL-227039).
- Asztalos, S., et al. (2007), "Properties of Ellipticity Correlation with Atmospheric Structure from Gemini South," *Astrophys. J.* **659** (1), 69–75 (UCRL-JRNL-226988).
- Bala, G., et al. (2007), "Combined Climate and Carbon-Cycle Effects of Large-Scale Deforestation," *P. Natl. Acad. Sci. USA* **104** (16), 6550–6555 (UCRL-JRNL-222088).
- Balhorn, R., et al. (2007), "Selective High-Affinity Ligand Antibody Mimics for Cancer Diagnosis and Therapy: Initial Application to Lymphoma/Leukemia," *Clin. Cancer Res.* **13** (18), 5621S–5628S.
- Ball, G. C., et al. (2007), "TRIUMF-ISAC Gamma-Ray Escape-Suppressed Spectrometer (TIGRESS): A Versatile Tool for Radioactive Beam Physics," *Nucl. Phys. A* **787**, 118C–125C.
- Banks, J. W., D. W. Schwendeman, A. K. Kapila, and W. D. Henshaw (2007), "A High-Resolution Godunov Method for Compressible Multimaterial Flow on Overlapping Grids," *J. Comput. Phys.* **223** (1), 262–297.
- Barton, N., et al. (2007), "Embedded Polycrystal Plasticity and Adaptive Sampling," *Int. J. Plasticity* **24** (2), 242–266 (UCRL-JRNL-226715).
- Bhaduri, D., S. K. Shukla, P. S. Graham, and M. B. Gokhale (2007), "Reliability Analysis of Large Circuits Using Scalable Techniques and Tools," *IEEE Trans. Circuit Syst-I* **54** (11), 2447–2460.
- Bono, M., et al. (2007), "Fabrication of Double Shell Targets with a Glass Inner Capsule Supported by SiO₂ Aerogel for Shots on the Omega Laser in 2006," *Fusion Sci. Technol.* **51** (4), 611–625.
- Bradley, D. K., et al. (2007), "Very-High-Growth-Factor Planar Ablative Rayleigh–Taylor Experiments," *Phys. Plasmas* **14** (5).
- Bronevetsky, G. and B. R. de Supinski (2007), "Complete Formal Specification of the OpenMP Memory Model," *Int. J. Parallel Prog.* **35** (4), 335–392 (UCRL-JRNL-226950).
- Chang, B. (2007), "A Deterministic Photon-Free Method to Solve Radiation Transfer Equations," *J. Comput. Phys.* **222** (1), 71–85.
- Chang, B. (2007), "The Conjugate Gradient Method Solves the Neutron Transport Equation h-optimally," *Numer. Linear Algebr.* **14** (10), 751–769.
- Chang, B. (2007), "The Incorporation of the Semi-Implicit Linear Equations into Newton's Method to Solve Radiation Transfer Equations," *J. Comput. Phys.* **226** (1), 852–878.
- Chang, B., et al. (2007), "Spatial Multigrid for Isotropic Neutron Transport," *SIAM J. Sci. Comput.* **29** (5), 1900–1917.
- Chen, G., Y. Liu, and N. Max (2007), "Real-Time View Synthesis from a Sparse Set of Views," *Signal Process.-Image* **22** (2), 188–202.
- Chen, Y., J. D. Cohen, and J. H. Krolik (2007), "Similarity-Guided Streamline Placement with Error Evaluation," *IEEE T. Vis. Comput. Gr.* **13** (6), 1448–1455.
- Choi, G. S., et al. (2007), "Performance Comparison of Coscheduling Algorithms for Non-Dedicated Clusters through a Generic Framework," *Int. J. High Perform. Comput.* **21** (1), 91–105.
- Choi, G. S., et al. (2007), "A Comprehensive Performance and Energy Consumption Analysis of Scheduling Alternatives in Clusters," *J. Supercomput.* **40** (2), 159–184.

- Chowell, G., T. Kostova, A. L. Rivas, and J. M. Hyman (2007), "Modeling the Spread of Foot-and-Mouth Disease," *CABI Reviews* (UCRL-JRNL-236739).
- Colella, P. and D. Trebotich (2007), "A Second-Order Accurate Conservative Front-Tracking Method in One Dimension," *SIAM J. Sci. Comput.* (UCRL-JRNL-235085).
- Crawford, T. D., et al. (2007), "PS13: An Open-Source Ab Initio Electronic Structure Package," *J. Comput. Chem.* **28** (9), 1610–1616.
- De Sterck, H., R. Falgout, J. Nolting, and U. M. Yang (2007), "Distance-Two Interpolation for Parallel Algebraic Multigrid," *Numer. Linear Algebr.* (UCRL-JRNL-230844).
- Dimits, A. M., et al. (2007), "Gyrokinetic Simulations of ETG and ITG Turbulence," *Nucl. Fusion* **47** (8), 817–824 (UCRL-JRNL-227861).
- Fattebert, J.-L., R. D. Hornung, and A. M. Wissink (2007), "Finite Element Approach for Density Functional Theory Calculations on Locally Refined Meshes," *J. Comput. Phys.* **223** (2), 759–773.
- Fisher, A., D. White, and G. Rodrigue (2007), "An Efficient Vector-Finite Element Method for Nonlinear Electromagnetic Modeling," *J. Comput. Phys.* **225** (2), 1331–1346 (UCRL-JRNL-222488).
- Garrett, P. E., et al. (2007), "Nuclear Structure Research at TRIUMF," *Acta Phys. Pol. B* **38** (4), 1169–1178.
- Glenzer, S. H., et al. (2007), "Experiments and Multiscale Simulations of Laser Propagation through Ignition-Scale Plasmas," *Nat. Phys.* **3** (10), 716–719.
- Griffith, B. E., R. D. Hornung, D. M. McQueen, and C. S. Peskin (2007), "An Adaptive, Formally Second-Order Accurate Version of the Immersed Boundary Method," *J. Comput. Phys.* **223** (1), 10–49.
- Gyulassy, A., V. Natarajan, V. Pascucci, and B. Hamann (2007), "Efficient Computation of Morse–Smale Complexes for Three-Dimensional Scalar Functions," *IEEE T. Vis. Comput. Gr.* **13** (6), 1440–1447.
- Gyulassy, A. G., et al. (2007), "Topologically Clean Distance Fields," *IEEE T. Vis. Comput. Gr.* **13** (6), 1432–1439.
- Homescu, C., L. Petzold, and R. Serban (2007), "Error Estimation for Reduced-Order Models of Dynamical Systems," *SIAM Rev.* **49** (2), 277–299.
- Hristova, K. R., et al. (2007), "Comparative Transcriptome Analysis of Methylibium Petroleiphilum PM1 Exposed to the Fuel Oxygenates Methyl Tert-Butyl Ether and Ethanol," *Appl. Environ. Microb.* **73** (22), 7347–7357.
- Ibarria, L., P. Lindstrom, and J. Rossignac (2007), "Spectral Interpolation on 3x3 Stencils for Prediction and Compression," *J. Computers* **2** (8), 53–63 (UCRL-JRNL-232201).
- Kane, S. R., et al. (2007), "Whole-Genome Analysis of the Methyl Tert-Butyl Ether-Degrading Beta-Proteobacterium Methylibium Petroleiphilum PM1," *J. Bacteriol.* **189** (5), 1931–1945 (UCRL-JRNL-226720).
- Kapila, A. K., D. W. Schwendeman, J. B. Bdzil, and W. D. Henshaw (2007), "Study of Detonation Diffraction in the Ignition-and-Growth Model," *Combust. Theor. Model.* **11** (5), 781–822.
- Kostova, T. (2007), "Persistence of Viral Infections on the Population Level Explained by an Immunoepidemiological Model," *Math. Biosci.* **206** (2), 309–319.
- Landen, O. L., et al. (2007), "The First Target Experiments on the National Ignition Facility," *Eur. Phys. J. D* **44** (2), 273–281.
- Lee, G., et al. (2007), "Dynamic Binary Instrumentation and Data Aggregation on Large-Scale Systems," *Int. J. Parallel Prog.* **35** (3), 207–232 (UCRL-JRNL-226801).
- Lee, S. L., and C. W. Gear (2007), "Second-Order Accurate Projective Integrators for Multiscale Problems," *J. Comput. Appl. Math.* **201** (1), 258–274.
- Li, Y., et al. (2007), "On the Origin of Smallpox: Correlating Variola Phylogenics with Historical Smallpox Records," *P. Natl. Acad. Sci. USA* **104** (40), 15787–15792.
- Liu, X., et al. (2007), "Uncertainties in Global Aerosol Simulations: Assessment Using Three Meteorological Datasets," *J. Geophys. Res.-Atmos.* **112** (UCRL-JRNL-215047).
- Liu, Y., et al. (2007), "Speech Recognition Systems on the Cell Broadband Engine Processor," *IBM J. Res. Dev.* **51** (5), 583–591.
- Marathe, J., et al. (2007), "METRIC: Memory Tracing via Dynamic Binary Rewriting to Identify Cache Inefficiencies," *ACM T. Progr. Lang. Sys.* **29** (2) (UCRL-JRNL-224540).
- Martz, H. E., et al. (2007), "Validation of Radiographic Simulation Codes Including X-Ray Phase Effects for Millimeter-Size Objects with Micrometer Structures," *J. Opt. Soc. Am. A* **24** (1), 169–178 (UCRL-JRNL-218242).
- Masters, N. D. and W. J. Ye (2007), "Octant Flux Splitting Information Preservation DSMC Method for Thermally Driven Flows," *J. Comput. Phys.* **226** (2), 2044–2062.
- McEnerney, J. and G. Stengle (2007), "Alternative Evidence for Non-Negativity," *J. Pure Appl. Algebra* **208** (3), 989–992.
- Miller, G. H. and D. Trebotich (2007), "Toward a Mesoscale Model for the Dynamics of Polymer Solutions," *J. Comput. Theor. Nanos.* **4** (4), 797–801.
- Mioduszewski, P. K., et al. (2007), "Power and Particle Handling and Wall Conditioning in NCSX," *Fusion Sci. Technol.* **51** (2), 238–260.

6.03 — Publications (cont.)

- Nilsson, S., N. A. Petersson, B. Sjogreen, and H.-O. Kreiss (2007), "Stable Difference Approximations for the Elastic Wave Equation in Second-Order Formulation," *SIAM J. Numer. Anal.* **45** (5), 1902–1936.
- Niski, K. and J. D. Cohen (2007), "Tile-Based Level of Detail for the Parallel Age," *IEEE T. Vis. Comput. Gr.* **13** (6), 1352–1359.
- Pascucci, V., G. Scorzelli, P.-T. Bremer, and A. Mascarenhas (2007), "Robust On-Line Computation of Reeb Graphs: Simplicity and Speed," *ACM T. Graphic.* **26** (3).
- Paul, B., et al. (2007), "Chromium Renderserver: Scalable and Open Remote Rendering Infrastructure," *IEEE T. Vis. Comput. Gr.*
- Reuland, M. et al. (2007), "Metal-Enriched Gaseous Halos around Distant Radio Galaxies: Clues to Feedback in Galaxy Formation," *Astron. J.* **133** (6), 2607–2623 (UCRL-JRNL-223328).
- Rodgers, A., et al. (2007), "Broadband Waveform Modeling of Moderate Earthquakes in the San Francisco Bay Area and Preliminary Assessment of the USGS 3D Seismic Velocity Model," *B. Seismol. Soc. Am.* (UCRL-228251).
- Serban, R., and C. Homescu, and L. Petzold (2007), "The Effect of Problem Perturbations on Nonlinear Dynamical Systems and Their Reduced-Order Models," *SIAM J. Sci. Comput.* **29** (6), 2621–2643.
- Singh, K., et al. (2007), "Predicting Parallel Application Performance via Machine Learning Approaches," *Concurr. Comp.-Pract. E.* **19** (17), 2219–2235 (UCRL-JRNL-222444).
- Sjogreen, B., K. Gustavsson, and R. Gudmundsson (2007), "A Model for Peak Formation in the Two-Phase Equations," *Math. Comput.* **76** (260), 1925–1940.
- Sjogreen, B. and N. A. Petersson (2007), "A Cartesian Embedded Boundary Method for Hyperbolic Conservation Laws," *Commun. Comput. Phys.* **2** (6), 1199–1219.
- Taylor, K. E., et al. (2007), "Estimating Shortwave Radiative Forcing and Response in Climate Models," *J. Climate* **20** (11), 2530–2543.
- Trebotich, D., G. H. Miller, and M. D. Bybee (2007), "A Hard Constraint Algorithm to Model Particle Interactions in DNA-LADEN Flows," *Nanoscale Microsc. Therm.* **11** (1–2), 121–128.
- Tripp, J. L., M. B. Gokhale, and K. D. Peterson (2007), "Trident: From High-Level Language to Hardware Circuitry," *Computer* **40** (3).
- Tsap, L. V., M. Duchaineau, D. B. Goldgof, and M. C. Shin (2007), "Data-Driven Feature Modeling, Recognition, and Analysis in a Discovery of Supersonic Cracks in Multimillion-Atom Simulations," *Pattern Recogn.* **40** (9), 2400–2407.
- Vassilevski, P. S. (2007), "Copper Mountain Special Issue on Iterative Methods," *SIAM J. Sci. Comput.* **29** (5) (UCRL-ABS-232680).
- Vo, H. T., et al. (2007), "Streaming Simplification of Tetrahedral Meshes," *IEEE T. Vis. Comput. Gr.* **13** (1), 145–155 (UCRL-JRNL-208710).
- Weber, G. H., P.-T. Bremer, and V. Pascucci (2007), "Topological Landscapes: A Terrain Metaphor for Scientific Data," *IEEE T. Vis. Comput. Gr.* **13** (6), 1416–1423.
- White, D. A., B. J. Fasenfest, R. N. Rieben, and M. L. Stowell (2007), "A QR Accelerated Volume-to-Surface Boundary Condition for the Finite-Element Solution of Eddy-Current Problems," *IEEE T. Magn.* **43** (5), 1920–1933.
- Xu, X. Q., et al. (2007), "Edge Gyrokinetic Theory and Continuum Simulations," *Nucl. Fusion* **47** (8), 809–816 (UCRL-JRNL-227862).
- Yee, H. C. and B. Sjogreen (2007), "Development of Low Dissipative High-Order Filter Schemes for Multiscale Navier–Stokes/MHD Systems," *J. Comput. Phys.*
- Yoon, S. E., and P. Lindstrom (2007), "Random-Accessible Compressed Triangle Meshes," *IEEE T. Vis. Comput. Gr.* **13** (6), 1536–1543.
- Zhou, C. E., et al. (2007), "MvirDB: A Microbial Database of Protein Toxins, Virulence Factors and Antibiotic Resistance Genes for Biodefense Applications," *Nucleic Acids Res.* **35**, D391–D394.

CONFERENCE PROCEEDINGS

- Amendt, P. et al. (2007), "Assessing the Prospects for Achieving Double-Shell Ignition on the National Ignition Facility Using Vacuum Hohlräume—Art. No. 056312," *Phys. Plasmas* (Amer. Inst. Physics), **14**, 56312–56312.
- Arnold, D., et al. (2007), "Stack Trace Analysis for Large-Scale Debugging," *IEEE/ACM International Parallel and Distributed Processing Symposium* (IEEE/ACM) (UCRL-CONF-227108).
- Balhorn, R., et al. (2007), "Selective High-Affinity Ligand Antibody Mimics for Cancer Diagnosis and Therapy: Initial Application to Lymphoma/Leukemia," *Clin. Cancer Res.* (Amer. Assoc. Cancer Research), **13**, 5621S–5628S.
- Ball, G. C., et al. (2007), "TRIUMF-ISAC Gamma-Ray Escape-Suppressed Spectrometer (TIGRESS): A Versatile Tool for Radioactive Beam Physics," *Nucl. Phys. A* (Elsevier Science BV), **787**, 118C–125C.
- Bono, M., et al. (2007), "Fabrication of Double-Shell Targets with a Glass Inner Capsule Supported by SiO₂ Aerogel for Shots on the Omega Laser in 2006," *Fusion Sci. Technol.* (Amer. Nuclear Society), **51**, 611–625.

- Bradley, D. K., et al. (2007), "Very-High-Growth-Factor Planar Ablative Rayleigh–Taylor Experiments—Art. No. 056313," *Phys. Plasmas* (Amer. Inst. Physics), **14**, 56313–56313.
- Bremer, P.-T., et al. (2007), "Topological Feature Extraction and Tracking," *J. Phys. Conf. Ser.*, **78**, 012007 (UCRL-CONF-232179), **78**, U76–U80.
- Bronevetsky, G. and B. R. de Supinski (2007), "Complete Formal Specification of the OpenMP Memory Model," *Int. J. Parallel Prog.* (Springer/Plenum Publishers), **35**, 335–392.
- Bronevetsky, G., et al. (2007), "Soft Error Vulnerability of Iterative Linear Algebra Methods," *The 2007 IEEE Workshop on Silicon Errors in Logic—System Effects (SELSE 3)* (IEEE) (UCRL-CONF-237305).
- Chen, Y., J. D. Cohen, and J. H. Krolik (2007), "Similarity-Guided Streamline Placement with Error Evaluation," *IEEE T. Vis. Comput. Gr.* (IEEE), **13**, 1448–1455.
- Childs, H. (2007), "Architectural Challenges and Solutions for Petascale Postprocessing," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78**, U118–U128 (UCRL-TR-232039).
- Curtis-Maury, M., et al. (2007), "Identifying Energy-Efficient Concurrency Levels Using Machine Learning," *International Workshop on Green Computing (GreenCom'07)* (IEEE) (UCRL-CONF-233024).
- Dahlgren, T. L. (2007), "Performance-Driven Interface Contract Enforcement for Scientific Components," *Lect. Notes Comput. Sc.* (Springer–Verlag Berlin), **4608**, 157–172 (UCRL-CONF-228332).
- De Sterck, H., et al. (2007), "Distance-Two Interpolation for Parallel Algebraic Multigrid," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78**, 012017 (UCRL-PROC-231994).
- de Supinski, B. R., et al. (2007), "Results of the PERI Survey of SciDAC Applications," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78** (UCRL-CONF-231869).
- Diachin, L., et al. (2007), "Interoperable Mesh and Geometry Tools for Advanced Petascale Simulations," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78**, U147–U152.
- Eder, D., et al. (2007), "Debris and Shrapnel Mitigation Procedure for NIF Experiments," *J. Phys. Conf. Ser.* (IOP Publishing, LTD) (UCRL-CONF-234490).
- Fisher, A. C., et al. (2007), "Hierarchical Material Models for Fragmentation Modeling in NIF–ALE–AMR," *J. Phys. Conf. Ser.* (IOP Publishing, LTD) (UCRL-CONF-234451).
- Gallagher, B. and T. Eliassi-Rad (2007), "An Examination of Experimental Methodology for Classifiers of Relational Data," *Proceedings of the Seventh IEEE International Conference on Data Mining* (IEEE) (UCRL-CONF-233643).
- Garrett, P. E., et al. (2007), "Nuclear Structure Research at TRIUMF," *Acta Phys. Pol. B* (Acta Physica Polonica B, Jagellonian Univ., Inst. Physics), **38**, 1169–1178.
- Gyulassy, A. G., et al. (2007), "Efficient Computation of Morse–Smale Complexes for Three-Dimensional Scalar Functions," *IEEE T. Vis. Comput. Gr.* (IEEE), **13**, 1440–1447.
- Gyulassy, A. G., et al. (2007), "Topologically Clean Distance Fields," *IEEE T. Vis. Comput. Gr.* (IEEE), **13**, 1432–1439.
- Ibarria, L., et al. (2007), "Spectral Predictors," *Data Compression Conference* (IEEE), 163–172 (UCRL-CONF-226261).
- Joy, K. I., et al. (2007), "Frameworks for Visualization at the Extreme Scale," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78**.
- Kegelmeyer, L. M., et al. (2007), "Local Area Signal-to-Noise Ratio (LASNR) Algorithm for Image Segmentation—Art. No. 66962H," *P. Soc. Photo-Opt. Instrum. Eng.* (SPIE—Int. Soc. Optical Engineering), **6696**, H6962–H6962.
- Koniges, A., et al. (2007), "Experiments for the Validation of Debris and Shrapnel Calculations," *J. Phys. Conf. Ser.* (IOP Publishing, LTD) (UCRL-CONF-234440).
- Koop, M., T. Jones, and D. K. Panda (2007), "Reducing Connection Memory Requirements of MPI for InfiniBand Clusters: A Message Coalescing Approach," *Proceedings of the Seventh IEEE International Symposium on Cluster Computing and the Grid Table of Contents* (ACM), 495–504 (UCRL-CONF-226304).
- Lee, B. C., et al. (2007), "Methods of Inference and Learning for Performance Modeling of Parallel Applications," *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (ACM) (UCRL-CONF-227097).
- Lee, G. L., et al. (2007), "Benchmarking the Stack Trace Analysis Tool for BlueGene/L," *International Conference on Parallel Computing (ParCo)* (UCRL-CONF-235241).
- Lee, G. L., et al. (2007), "Pynamic: the Python Dynamic Benchmark," *IEEE International Symposium on Workload Characterization* (IEEE), 101–106 (UCRL-CONF-232621).
- Love, N. S. and C. Kamath (2007), "Image Analysis for the Identification of Coherent Structures in Plasma—Art. No. 66960D," *P. Soc. Photo-Opt. Instrum. Eng.* (SPIE—Int. Soc. Optical Engineering), **6696**, D6960–D6960.
- McInnes, L. C., et al. (2007), "Research Initiatives for Plug-and-Play Scientific Computing," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78** (UCRL-CONF-232026).
- Miller, G. H. and D. Trebotich (2007), "Toward a Mesoscale Model for the Dynamics of Polymer Solutions," *J. Comput. Theor. Nanos.* (Amer. Scientific Publishers), **4**, 797–801.
- Mirin, A. and P. Worley (2007), "Extending Scalability of the Community Atmosphere Model," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78**, U610–U616 (UCRL-PROC-232512).

6.03 — Publications (cont.)

- Niski, K. and J. D. Cohen (2007), "Tile-Based Level of Detail for the Parallel Age," *IEEE T. Vis. Comput. Gr.* (IEEE), **13**, 1352–1359.
- Noeth, M., et al. (2007), "Scalable Compression and Replay of Communication Traces in Massively Parallel Environments," *21st International Parallel and Distributed Processing Symposium (IPDPS 2007)* (IEEE) (UCRL-CONF-227098).
- Panas, T., et al. (2007), "Communicating Software Architecture Using a Unified Single-View Visualization," (IEEE) 217–226.
- Pascucci, V., et al. (2007), "Robust Online Computation of Reeb Graphs: Simplicity and Speed—Art. No. 58," *ACM T. Graphic.* (ACM), **26**, 58.
- Pon, R. K., et al. (2007), "iScore: Measuring the Interestingness of Articles in a Limited User Environment," *IEEE Symposium on Computational Intelligence and Data Mining* (IEEE), 354–361 (UCRL-CONF-225289).
- Pon, R. K., et al. (2007), "Tracking Multiple Topics for Finding Interesting Articles," *Proceedings of the Seventh IEEE International Symposium on Cluster Computing and the Grid Table of Contents* (ACM), 560–569 (UCRL-CONF-228286).
- Pope, G. (2007), "Measuring the 'Good' in Good Enough Software," *STARWest 2007 Conference Proceedings* (Software Quality Engineering) (UCRL-PRES-223524).
- Rountree, B., et al. (2007), "Bounding Energy Consumption in Large-Scale MPI Programs," *Supercomputing* (IEEE/ACM) (UCRL-CONF-233221).
- Ruppert, S. D., et al. (2007), "Enhancing Seismic Calibration Research through Software Automation and Scientific Information Management," (DOE/DOD), **1**, 965–974 (UCRL-CONF-232583).
- Schulz, M. and B. R. de Supinski (2007), "P^NMPI Tools: A Whole Lot Greater than the Sum of Their Parts," *Supercomputing* (IEEE/ACM) (UCRL-CONF-229978).
- Schulz, M. and B. R. de Supinski (2007), "Practical Differential Profiling," *Lect. Notes Comput. Sci.* (Springer-Verlag Berlin), **4641**, 97–106 (UCRL-CONF-227812).
- Tong, H., et al. (2007), "Fast Best-Effort Pattern Matching in Large Attributed Graphs," *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM), **978-1-59593-609-7**, 737–746 (UCRL-CONF-231426).
- Trebotich, D. (2007), "Simulation of Biological Flow and Transport in Complex Geometries Using Embedded Boundary Volume-of-Fluid Methods," *J. Phys. Conf. Ser.* (IOP Publishing, LTD), **78**, U566–U570.
- Trinitis, C. and M. Schulz (2007), "Sixth International Special Session on: Current Trends in Numerical Simulation for Parallel Engineering Environments ParSim 2005," *Lect. Notes Comput. Sci.* (Springer-Verlag), **4757**, 354–355 (UCRL-PROC-232591).
- Weber, G. H., P.-T. Bremer, and V. Pascucci (2007), "Topological Landscapes: A Terrain Metaphor for Scientific Data," *IEEE T. Vis. Comput. Gr.* (IEEE), **13**, 1416–1423.
- Weber, G. H., et al. (2007), "Visualization Tools for Adaptive Mesh Refinement Data," *Proceedings of the 4th High-End Visualization Workshop* (Lehmanns Media), **4**, 12–25.
- Yee, H. C. and B. Sjogreen (2007), "Development of Low Dissipative High-Order Filter Schemes for Multiscale Navier–Stokes/MHD Systems," *Astron. Soc. Pacific CS* (Academic Press Professional), **47**.

6.04 — Acronyms and Abbreviations

2D – two-dimensional	ITIL – Information Technology Infrastructure Library
3D – three-dimensional	ITPD – Information Technology Protection Division
BNT – Backbone Network Team	LC – Livermore Computing
ADEPT – Application Development Environment and Performance Team	LCRM – Livermore Computing Resource Management
AFP – Argus Field Processor	LLNL – Lawrence Livermore National Laboratory
ALE – Arbitrary Lagrangian–Eulerian	LLNS – Lawrence Livermore National Security, LLC
AMR – adaptive mesh refinement	LPOM – Laser Performance Operations Model
ASC – Advanced Simulation and Computing	LUSciD – Livermore–UC San Diego Scientific Data Study
CANES – Context-Aware Nuclear Evaluation System	MC – mediated component
CIVS – Chamber Interior Viewing System	MHD – magnetohydrodynamics
CMELS – Chemistry, Materials, Earth, and Life Sciences	MPI – Message Passing Interface
CMIP3 – Climate Model Intercomparison Project, Phase 3	NDR – NIF Data Repository
CMS – component mediation system	NIF – National Ignition Facility
CORBA – Common Object Request Broker Architecture	NNSA – National Nuclear Security Administration
CRI – Cluster Resources, Inc.	PCMDI – Program for Climate Model Diagnosis and Intercomparison
DHS – Department of Homeland Security	RIID – radiation isotope identification devices
DoD – Department of Defense	RPM – radiation portal monitor
DOE – Department of Energy	SMSG – System Management Solution Group
DPM – data pedigree management	SPU – special processing unit
DS – double-shell	SQE – software quality engineering
HPC – high-performance computing	STAT – Stack Trace Analysis Tool
ICCS – Integrated Computer Control System	TF – trillion floating-point operations per second
ICF – inertial confinement fusion	TLCC – Tri-Laboratory Capacity Cluster
ICRMG – Integrated Computational Resource Management Group	TOSS – Tripod Operating System Software
IPCC – Intergovernmental Panel on Climate Change	V&V – verification and validation
ISCR – Institute for Scientific Computing Research	VDP – Viral Detection Project
IT – information technology	WCI – Weapons and Complex Integration

